

FP²: Fully in-Place Functional Programming

Technical Report MSR-TR-2023-19, 2023-07-03, v2

ANTON LORENZEN, University of Edinburgh, UK

DAAN LEIJEN, Microsoft Research, USA

WOUTER SWIERSTRA, Universiteit Utrecht, Netherlands

As functional programmers we always face a dilemma: should we write purely functional code, or sacrifice purity for efficiency and resort to in-place updates? This paper identifies precisely when we can have the best of both worlds: a wide class of purely functional programs can be executed safely using in-place updates without requiring allocation, provided their arguments are not shared elsewhere. We describe a linear *fully in-place* (FIP) calculus where we prove that we can always execute such functions in a way that requires no (de)allocation and uses constant stack space. Of course, such a calculus is only relevant if we can express interesting algorithms; we provide numerous examples of in-place functions on datastructures such as splay trees or finger trees, together with in-place versions of merge sort and quick sort. We also show how we can generically derive a map function over *any* polynomial data type that is fully in-place. Finally, we have implemented the rules of the FIP calculus in the Koka language. Using the Perceus reference counting garbage collection, this implementation dynamically executes FIP functions in-place whenever possible.

Additional Key Words and Phrases: FBIP, Tail Recursion Modulo Cons

1 INTRODUCTION AND OVERVIEW

The functional program for reversing a list in linear time using an accumulating parameter has been known for decades, dating back at least as far as Hughes’s work on difference lists [1986]:

```
fun reverse-acc( xs : list<a>, acc : list<a> ) : list<a>
  match xs
  Cons(x,xx) -> reverse-acc( xx, Cons(x,acc) )
  Nil       -> acc

fun reverse( xs : list<a> ) : list<a>
  reverse-acc(xs,Nil)
```

As this definition is *pure*, we can calculate with it using equational reasoning in the style of Bird and Meertens [Backhouse 1988; Gibbons 1994]. Using simple induction, we can, for instance, prove that this linear time list reversal produces the same results as its naive quadratic counterpart.

Not all in the garden is rosy: what about the function’s memory usage? The purely functional definition of `reverse` allocates fresh `Cons` nodes in each iteration; an automatic garbage collector needs to discard unused memory. This generally induces a performance penalty relative to an imperative in-place implementation that destructively updates the pointers of a linked list. Reasoning about such imperative in-place algorithms, however, is much more difficult.

As programmers we seem to face a dilemma: should we write purely functional code, or sacrifice purity for efficiency and resort to in-place updates? This paper identifies precisely when we can have the best of both worlds: a wide class of purely functional programs can be executed safely using in-place updates without requiring allocation, including the `reverse` function above.

In particular, what *if* the compiler can make the assumption that the function parameters are *owned* and unique at runtime, i.e. that there are no other references to the input list `xs` of `reverse` at runtime. In that case, the compiler can safely *reuse* any matched `Cons` node and update it in-place with the result – effectively updating the list in-place. In this paper we describe a novel *fully in-place* (FIP) calculus that guarantees that such a function can be compiled in a way to never (de)allocate memory or use unbounded stack space—it can be executed fully in-place.

To illustrate the purely functional fully in-place paradigm, we consider *splay trees* as described by Sleator and Tarjan [1985]. These are self-balancing trees where every access to an element in the tree, including lookup, restructures the tree such that the element is “splayed” to the top of the tree. As a result, the `lookup` function not only returns a boolean representing whether or not the element was found in the tree, but also a newly splayed tree. Splay trees are generally not considered well-suited for functional languages, because every such restructuring of the tree copies the spine of the tree leading to decreased performance relative to an imperative implementation that can rebalance the tree in-place. Surprisingly, it turns out to be possible to write the splay algorithms in a purely functional style using *fully in-place* functions.

1.1 Zippers and Unboxed Tuples

Let us first define the type of splay trees containing integers¹:

```
type stree
  Node(left : stree, value : int, right : stree)
  Leaf
```

For the `lookup` function, once we find a given element in the tree, we need to somehow navigate back through the tree to splay the node to the top. The usual imperative approach uses parent pointers for this, but in a purely functional style we can use Huet’s *zipper* [1997] instead. The central idea is a simple one: to navigate through a tree step by step, we store the current subtree in focus together with its context. Naively, one might try to represent the context as a path from the root of the tree to the current subtree. The zipper, however, *reverses* this path so each step up or down the tree requires only constant time. For splay trees, we can define the corresponding zipper as:

```
type zipper
  Root
  NodeL( up : zipper, value : int, right : stree )
  NodeR( left : stree, value : int, up : zipper )
```

Huet already observed that the zipper operations could be implemented in-place. The original paper presenting the zipper concludes with the following paragraph [Huet 1997]:

Efficient destructive algorithms on binary trees may be programmed with these completely applicative primitives, which all use constant time, since they all reduce to local pointer manipulation.

Our FIP calculus provides the language to make such a statement precise: using the rules presented in the next section we can check statically that the various operations on zippers are indeed fully in-place. For example, we can move focus to the left subtree as follows:

```
fip fun left(t : stree, ctx : zipper) : (stree, zipper)
  match! t
    Node(l,x,r) -> (l, NodeL(ctx,x,r))
    Leaf       -> (Leaf, ctx)
```

The `fip` keyword indicates that a static check guarantees the function is fully in-place. This check, specified in Section 2, verifies that the function lives in a *linear* fragment of the language where the function parameters and variables are *owned* and can only be used once. As a consequence, we can safely reuse the memory of a linear value in-place once it is no longer used. The `match!` keyword² indicates a *destructive match*, after which the matched variable `t` can no longer be used. Intuitively, we can then see straightaway that the matched `Node` cell has become redundant and can be reused

¹All examples in this paper are written in the Koka language which has a full implementation of the FIP check (v2.4.2).

²In our implementation it turns out we can always *infer* when a match needs to be destructive and we can always write just `match` for both destructive and borrowing matches. However, for clarity and correspondence to our formal FIP calculus, we denote destructive matches explicitly in this paper.




Fig. 1. Looking up the number 3 in a splay tree (A), creating the zipper context to the node containing 3 (B), and splaying the node up to the top (C).

for the `NodeL` cell of the zipper.

To make this intuition more precise, we view a destructively matched constructor, such as `Node(1,x,r)`, as a collection of its children `1`, `x`, and `r`, and a *reuse credit* \diamond_3 . This “diamond” resource type \diamond_k represents a specific heap cell of size k and is inspired by the work of Aspinall and Hofmann on space credits [Aspinall and Hofmann 2002; Aspinall et al. 2008; Hofmann 2000b 2000a]. Similar to their space credits, we also apply the linearity restriction to reuse credits, but, unlike space credits, a reuse credit can not be split or merged with other credits. In our example, the `match!` introduces a reuse credit \diamond_3 for the `Node`, which is consumed by the allocation of `NodeL` which allows our `left` function to be fully in-place.

It may seem that we still need to allocate a tuple to store the result. Such allocation, however, is usually unnecessary since tuples are often created only to hold multiple return values and immediately destructed afterwards. In our calculus and implementation, we model tuples as *unboxed* values hence no allocation is needed for these³.

We can now also see that the list reversal of the introduction is also fully in-place:

```
fip fun reverse-acc( xs : list<a>, acc : list<a> ) : list<a>
  match! xs
  Cons(x,xx) -> reverse-acc( xx, Cons(x,acc) )
  Nil       -> acc
```

where the destructive match on `xs` allows the matched `Cons` cell to be reused (\diamond_2) for the `Cons(x,acc)` allocation in the branch.

1.2 Splay Tree Lookup and Atoms

Using the zipper definition for splay trees, we can now define the `lookup` function as follows:

```
fip fun lookup( t : stree, x : int ) : (bool, stree)
  zlookup(t,x,Root)

fip fun zlookup( t : stree, x : int, ctx : zipper ) : (bool, stree)
  match! t
  Leaf -> (False, splay-leaf(ctx)) // not found, but splay anyway
  Node(1,y,r) ->
    if x < y then zlookup(1,x,NodeL(ctx,y,r)) // go down the left (NodeL reuses Node)
    elif x > y then zlookup(r,x,NodeR(1,y,ctx)) // go down the right (NodeR reuses Node)
    else (True, splay(Top(1,y,r),ctx)) // found it, now splay it to the top
```

The `lookup` function calls `zlookup` with an initial empty context `Root`. It seems we need to allocate the `Root` constructor, but constructors without any fields do not require allocation. These are typically implemented using pointer-tagging where only the tag is used to represent them. We call such constructors *atoms*. Examples include the `Nil` of lists, but also booleans (`True` and `False`), and, depending on the implementation, primitive types like integers (`int`) or floats.

³In Koka, tuples are implemented as *value types* which are unboxed and passed in registers. The `fip` keyword additionally checks that no automatic (heap allocated) boxing is applied for such value types inside a FIP function.

FP²: Fully in-Place Functional Programming

The `zlookup` function traverses down the tree while extending the current zipper context *in-place* with the path that is followed. Figure 1 shows a concrete example of how `zlookup` constructs the zipper in the transition from (A) to (B). Once the element is found, the corresponding node is splayed back up to the top of the tree as `splay(Top(l,x,r),ctx)`. Here we use the `Top` constructor:

```
type top
  Top( left : stree, value : int, right : stree )
```

But why is this `Top` constructor needed? Can we not just call `splay` directly with explicit arguments as `splay(l,x,r,ctx)`? This is not possible though in a fully in-place way. In particular, it would mean that the `Node(l,y,r)` on which we matched would need to be deallocated as it cannot be reused immediately (and deallocation is not allowed in `fib` functions). Moreover, we would now need to allocate the final top node later on. If we define `splay` without using `Top`, we would get something like:

```
fun splay( l : stree, x : int, r : stree, ctx : szipper ) : stree // not fib!
  match! ctx
  Root -> Node(l,x,r)
  ...
```

That is, once the zipper `ctx` is at the `Root` we need to return a fresh `Node` with `x` on top. As there is no constructor that can be reused (since `Root` is just an atom), this would require allocation. By using the intermediate `Top` constructor we avoid this: at the call site we can now reuse the `Node` that we just matched for `Top`, and later on when we reach the `Root`, we can reuse `Top` again to create the final `Node` that becomes the top of the returned `splay` tree. This results in the final fully in-place definition of the `splay` function:

```
fib fun splay( top : top , ctx : szipper ) : stree
  match! top
  Top(l,x,r) -> match! ctx
    Root -> Node(l,x,r)
    NodeL(Root,y,ry) -> Node(l,x,Node(r,y,ry)) // zig
    NodeL(NodeR(lz,z,up),y,ry) -> splay( Top(Node(lz,z,l),x,Node(r,y,ry)), up) // zig-zag
    NodeL(NodeL(up,z,rz),y,ry) -> splay( Top(l,x,Node(r,y,Node(ry,z,rz))), up) // zig-zig
    NodeR(l,y,y,Root) -> Node(Node(l,y,y,l),x,r)
    NodeR(l,y,y,NodeL(up,z,rz)) -> splay( Top(Node(l,y,y,l),x,Node(r,z,rz)), up) // (B)->(C)
    NodeR(l,y,y,NodeR(lz,z,up)) -> splay( Top(Node(Node(lz,z,ly),y,l),x,r), up)
```

The compiler statically checks that this function is fully in-place. As a result, we know that its execution uses no stack space and performs all its rebalancing operations without any (de)allocation – each `Top` and `Node` can reuse a destructively matched `Top`, `NodeL`, or `NodeR` in every branch. Each of the matched cases correspond to a “zig”, “zig-zig”, and “zig-zag” rebalance operation as described by Sleator and Tarjan [1985]. Figure 1 shows a concrete example of the “zig-zag” in the transition from (B) to (C). For completeness, the auxiliary `splay-leaf` function that is called in case the item is not a member of the tree is included below:

```
fib fun splay-leaf( ctx : szipper ) : stree
  match! ctx
  Root -> Leaf
  NodeL(up,x,r) -> splay(Top(Leaf,x,r),up)
  NodeR(l,x,up) -> splay(Top(l,x,Leaf),up)
```

The definition of `splay` may look somewhat involved but compared to the imperative definition it is fairly concise. Moreover, the usual imperative algorithm uses extra space for parent pointers in each node. We do not need this: if we study the generated code for the `fib` functions, we see that the reuse of the zipper nodes corresponds to the usual “pointer-reversal” techniques [Schorr and Waite 1967] (which is illustrated nicely in Figure 1 (B)). Such pointer-reversal is not often used explicitly in practice though since it is difficult to get right by hand. In the code above, however, the fully in-place `fib` functions using zippers provide a statically typed, memory safe, purely functional definition with the same runtime behaviour, without requiring explicit pointer manipulation.

1.3 Borrowing and Second-Order Functions

While our examples have so far been entirely first-order, we also allow functions to be passed as arguments. For example, we can map over a splay tree as follows:

```
fbip fun smap( t : stree, ^f : int -> int ) : stree
  match! t
    Node(l,v,r) -> Node(smap(f,l), f(v), smap(f,r))
    Leaf       -> Leaf
```

In this function we seemingly violate our linearity constraint since f is used three times in the first branch. However, the function parameter is marked as *borrowed* using the hat notation (f) [Ulrich and de Moura 2019]. This allows the parameter to be freely shared but at the same time it cannot be used in a destructive match, passed as an owned parameter, or returned as a result. Such borrowing is often useful for functions that inspect a data structure. Consider the following example `is-node` function:

```
fip is-node( ^t : stree ) : bool
  match t
    Node(.,.,_) -> True
    _          -> False
```

This function would not be fully in-place if we matched destructively. A subtle point about higher-order functions is that we consider an application $f(e)$ as a borrowed use of f , and, as a consequence, f cannot modify any captured free variables in-place. We enforce this in our calculus by only allowing top-level functions (rather than arbitrary closures) as arguments in our fully in-place calculus, effectively making it second-order.

Finally, note that the `smap` function is marked not as `fip` but as `fbip`. Unlike the earlier `splay` function, `smap` has recursive calls in non-tail positions. That makes it hard to claim that this function is “fully in-place”—after all, its execution uses stack space linear in the depth of the splay tree. We use the `fbip` keyword to signify FIP functions that still reuse in-place but are allowed to use arbitrary stack space and deallocate memory. Nevertheless, for `smap` this is not really required: in Section 3 we show that *any* map function of polynomial datatypes (including `smap`) can be transformed into a tail-recursive, zipper-based traversal that *is* fully in-place.

The `fbip` keyword is derived from the “functional but in-place” technique [Reinking, Xie et al. 2021] which is a more liberal notion of our strict fully in-place functions. Our implementation also supports `fip(n)` and `fbip(n)` for a constant n , which allows the function to allocate at most a constant n constructors. This is sometimes useful for functions like splay tree insertion where a single `Node` may need to be allocated for the newly inserted element, making it `fip(1)`.

1.4 Fully in-Place in a Functional World

One might argue that fully in-place programming is just imperative programming in functional clothing. Where are the closures, the non-linear values, the persistence? And who allocates the list to be reversed in the first place? We need to be able to embed our fully in-place functions in a larger host language to be useful. The challenge is to do this safely while still guaranteeing in-place updates when possible. To illustrate this point, consider the following function:

```
fun palindrome( xs : list<a> ) : list<a>
  append(xs, reverse(xs))
```

Even though `reverse` is a `fip` function, it would not be safe for it to destructively update its input list since the argument `xs` is used twice (as an owned argument) in the body of `palindrome`. The FIP calculus presented here statically checks a function’s definition—yet deciding which *calls* to `fip` functions can be safely executed using destructive updates requires further information about how arguments are shared at call sites.

FP²: Fully in-Place Functional Programming

1.4.1 Uniqueness Typing. One way to check this information statically is using using a *uniqueness* type system. For example, Clean [Brus et al. 1987] is a functional language where the type system tracks when arguments are *unique* or *shared* [Barendsen and Smetsers 1995; De Vries et al. 2008]. A `!fip` function may safely use in-place mutation, provided all owned parameters have a *unique type*. In that way, the type system guarantees that any argument passed at runtime will be a unique reference to that object, ruling out any possible sharing. As a result, it is always safe to reuse the argument in-place. One possible drawback of linear type systems and uniqueness typing is that it leads to code duplication, where a single function can have multiple different implementations: one version taking a unique argument; and one taking a shared argument. For example, we may need to define two `reverse` functions that have an equivalent implementation but only differ in the `!fip` annotation and the uniqueness type of the input list – one uses copying and can be used persistently with a shared list, while the FIP variant updates the list in-place but requires the input list to be unique.

1.4.2 Precise Reference Counting. Checking call sites of `!fip` functions need not happen statically. Instead, we could also use a *dynamic* approach where we check at runtime if a FIP function can be executed in-place [Lorenzen and Leijen 2022; Schulte and Grieskamp 1992; Ullrich and de Moura 2019]. This is the approach taken in our implementation in the Koka language, which uses Perceus precise reference counting [Reinking, Xie et al. 2021; Ullrich and de Moura 2019]. When an object has a reference count of one, it is safe to update it in-place. To illustrate how the compiled code looks in practice, consider the compilation of the fully in-place `reverse-acc` function. The destructive match on the input list now dynamically checks whether or not the list can be mutated in place and the generated code will look something like:

```
fip fun reverse-acc( xs : list<a>, acc : list<a> ) : list<a>
  match! xs
  Cons(x,xx) ->
    val ru = if is-unique(xs) then &xs else { dup(x); dup(xx); decref(xs); alloc(2) }
    reverse-acc( xx, Cons@ru(x,acc) )
  Nil -> acc
```

The reuse credit \diamond_2 is compiled into an explicitly named reuse token `ru`, and holds to the memory location of the resulting `Cons` cell. If the input list is unique, we *reuse* the address of the input, `&xs`, and otherwise, we adjust the reference counts of the children accordingly and return freshly allocated memory of the right size. In the recursive call, we initialize the `Cons` cell in-place at the `ru` memory location as `Cons@ru(x,acc)`.

Compared to the static analysis, we have lost the static guarantee that the owned parameters are unique at runtime, but also we gained expressiveness: in particular, we can define now a *single* purely functional but fully in-place `reverse` function that serves all usage scenarios: it efficiently updates the elements in place if the argument list is unique at runtime, but it also adapts dynamically when the list, or any sublist happens to be shared – falling back gracefully to allocating fresh `Cons` nodes for the resulting list.

We show formally that the FIP calculus can be embedded, as a subcalculus, into a Perceus-style linear resource calculus λ^{fip} with unboxing, atoms and borrowing in Section 5.

This language can now express the full range of functional programming. Yet, only one rule needs to be extended: rather than always reusing in a destructive `match!`, we inspect the reference counts and only reuse unique cells as above. In particular, this shows that fully in-place functions can execute on unique arguments with no overhead due to reference counting (besides the `is-unique` check).

In this setting, when we call a function like `map` fully in-place, it means that the function itself does not intrinsically allocate any memory. However, any particular evaluation can still allocate by passing in a non-`!fip` function, for example `map(xs, fn(x) Cons(x,Nil))`. That is an essential part

of our approach though, where we are able to reason about a fully in-place function like `map` in isolation but still allow it to be used in allocating scenarios from a non-FIP context.

In the next section we look at a precise definition of a fully in-place calculus.

1.5 Contributions

To support the motivating examples outlined so far, this paper makes the following contributions:

- Following the pioneering work on the LFPL calculus [Hofmann 2000b 2000a], we present a novel *fully in-place* (FIP) calculus (Section 2), precisely capturing those functions that can be executed fully in-place. We provide a standard functional operational semantics for our language, but also define an equivalent semantics for FIP functions in terms of a fixed store, where no (de)allocation can take place. As a result, we know that FIP functions never allocate memory and use bounded stack space. As shown for splay trees, *atoms* and *unboxed tuples* are needed to avoid allocations for many common scenarios and the FIP calculus includes these features. Furthermore, the rules of the FIP calculus provide a static guarantee of linearity in a *syntactic* way where parameters can be *owned* uniquely or *borrowed*.
- The FIP calculus is only useful if it can actually be used to describe interesting algorithms. To show the wide applicability of our approach we present a variety of familiar functional programs and operations on datastructures that are all *fully in-place*. We have already seen how Huet’s *zipper* [1997] datastructure, colloquially described as the functional equivalent of backpointers, can be used fully in-place, and how we can use this to implement fully in-place *splay trees* [Sleator and Tarjan 1985]. In Section 3, we further show that we can use a defunctionalized CPS transformation [Danvy 2008] to derive a generic `map` function for *any* polynomial inductive datatype. The derived `map` uses fully in-place Schorr-Waite traversal [Schorr and Waite 1967] without using any extra stack- or heap space. In Section 4, we show further examples of functional algorithms and datastructures that are fully in-place, including imperative red-black tree insertion [Cormen et al. 2022], `cons` and `append` operations on finger trees [Claessen 2020; Hinze and Paterson 2006], and even sorting algorithms like merge sort and quick sort. We have an implementation of the FIP calculus in a fork of the Koka compiler that can check and compile all examples in this paper.
- Finally, we study the dynamic embedding of our FIP calculus based on precise reference counting in detail in Section 5. Integrating the static FIP calculus with the dynamic Perceus linear resource calculus, λ^1 [Lorenzen and Leijen 2022; Reinking, Xie et al. 2021], gives us precisely the information we need to decide when a function call can be executed in-place or not. However, the original linear resource calculus does not model reuse, atoms, unboxing, or borrowing, all of which are essential for FIP programs. We simplify and extend the linear resource calculus into a new calculus (λ^{fip}) which includes all these features, and give a novel proof of soundness. Surprisingly, it turns out that the extended linear resource calculus λ^{fip} can be seen a pure *extension* of the FIP calculus; and the rules of the FIP calculus are a strict subset of λ^{fip} . In particular, FIP is exactly that subset of λ^{fip} which requires no dynamic reference counting or memory management at runtime. As a result, in the Perceus setting FIP functions can interact safely with any other function, executing in-place when possible and copying when necessary.
- In Section 6 we show a short performance evaluation of our particular implementation. It shows that `fip` algorithms are competitive to standard functional algorithms in Koka. This is somewhat expected since the standard algorithms can already avoid many allocations through the existing dynamic reuse as part of Perceus reference counting. If such dynamic reuse is disabled for the standard algorithms, `fip` functions tend to outperform with a larger margin.

2 A LANGUAGE FOR FULLY IN-PLACE UPDATE

Figure 2 presents the syntax of the *fully in-place* FIP calculus. The syntax has been carefully chosen to be expressive enough to cover many interesting functions as shown in this paper, but at the same

FP²: Fully in-Place Functional Programming

Expressions:

$e ::= (v, \dots, v)$	(unboxed tuple)	$v ::= x, y$	(variables)
$e e$	(application)	$C^k v_1 \dots v_k$	(constructor of arity k)
$f(e; e)$	(call)		
$\text{let } \bar{x} = e \text{ in } e$	(let binding)		
$\text{match } e \{ \bar{p} \mapsto e \}$	(matching)	$p ::= C^k x_1 \dots x_k$	(pattern)
$\text{match! } e \{ \bar{p} \mapsto e \}$	(destructive match)		

$\Sigma ::= \emptyset \mid \Sigma, f(\bar{y}; \bar{x}) = e$ (recursive top-level functions with borrowed parameters \bar{y})

Syntax:

$\bar{v} \doteq (v_1, \dots, v_k)$	$(k \geq 1)$	$\text{let } x = e_1 \text{ in } e_2 \doteq$	$\text{let } (x) = e_1 \text{ in } e_2$
$\bar{x} \doteq (x_1, \dots, x_k)$	$(k \geq 1)$	$\lambda x_1, \dots, x_k. e \doteq$	$\lambda (x_1, \dots, x_k). e$
$v \doteq (v)$	(unboxed singleton)		

Fig. 2. Syntax of the FIP calculus.

time restricted enough to be straightforward to analyze. Particular properties of our syntax are the inclusion of unboxed tuples, borrowed parameters, and the lack of general lambda expressions.

The syntax distinguishes between expressions e , and values v that cannot be evaluated further. Values are either variables or fully applied constructors C^k , taking k values as arguments. We often leave out the superscript k when not needed.

Unboxed tuples (v_1, \dots, v_k) are considered expressions, rather than values. In this way, we syntactically rule out that unboxed tuples may be passed as an argument to a constructor, causing them to become “boxed” (and allocated). Instead of enforcing this with a type system [Peyton Jones and Launchbury 1991], we use this syntactic restriction to enforce this property. By doing so, the check is simpler and allows us to specify the FIP calculus independent of its static semantics.

We often write just v or e for a singleton unboxed tuple (v) , and write an overline to denote an unboxed tuple (v_1, \dots, v_k) as \bar{v} , or an unboxed tuple of variables (x_1, \dots, x_k) as \bar{x} . Since expressions always eventually evaluate to an unboxed tuple, the $\text{let } \bar{x} = e_1 \text{ in } e_2$ expression binds all components of the result unboxed tuple e_1 in \bar{x} .

There are no general lambda expressions. In general, closures need to be heap allocated if they contain free variables. To keep the FIP rules as simple as possible, we do not allow *arbitrary* lambda expressions. Instead, the global Σ environment holds all top-level functions f , which can be mutually recursive and passed as arguments. This makes our calculus essentially second-order. A top-level function is declared as $f(\bar{y}; \bar{x}) = e$, where \bar{y} are the *borrowed* parameters, and \bar{x} are the *owned* (unique) parameters. Just as in a let binding, the \bar{y} and \bar{x} bind the components of the unboxed tuples that are passed. A function is called by writing $f(e_1; e_2)$ with e_1 for the borrowed arguments, and e_2 for the owned arguments.

If there are no borrowed arguments, we sometimes write just $f(e_2)$. The syntax $e_1 e_2$ is used for general application when the function to be called is not statically known. This happens when a function f is passed as an second-order argument itself, and in such case, e_2 is always passed as the owned parameter(s). We have two match expressions, the regular match, and the destructive match!. There is no difference in the functional semantics between the two, but in a heap semantics the destructive match can be used for reuse, and as we see in the FIP rules in Figure 4, it can only be used on owned parameters.

2.1 Functional Operational Semantics

Figure 3 gives the functional operational semantics for our calculus. This semantics does not yet use a heap. An evaluation context E is a term with a single occurrence of a hole \square in place of a

Evaluation order:

$$E ::= \square \mid E e \mid \bar{v} E \mid f(E; e) \mid f(\bar{v}; E) \mid \text{let } \bar{x} = E \text{ in } e \quad \frac{e_1 \longrightarrow e_2}{E[e_1] \mapsto E[e_2]} \text{ STEP}$$

$$\mid \text{match } E \{ \bar{p} \mapsto \bar{e} \} \mid \text{match! } E \{ \bar{p} \mapsto \bar{e} \}$$

Evaluation steps:

$$\begin{array}{llll} (\text{let}) & \text{let } \bar{x} = \bar{v} \text{ in } e & \longrightarrow & e[\bar{x}:=\bar{v}] \\ (\text{call}) & f(\bar{v}_1; \bar{v}_2) & \longrightarrow & e[\bar{y}:=\bar{v}_1, \bar{x}:=\bar{v}_2] \quad \text{with } f(\bar{y}; \bar{x}) = e \in \Sigma \\ (\text{app}) & (f) \bar{v} & \longrightarrow & e[\bar{x}:=\bar{v}] \quad \text{with } f(; \bar{x}) = e \in \Sigma \\ (\text{match}) & \text{match } (C \bar{v}) \{ \bar{p} \mapsto \bar{e} \} & \longrightarrow & e_i[\bar{y}:=\bar{v}] \quad \text{with } p_i = C \bar{y} \\ (\text{match!}) & \text{match! } (C \bar{v}) \{ \bar{p} \mapsto \bar{e} \} & \longrightarrow & e_i[\bar{x}:=\bar{v}] \quad \text{with } p_i = C \bar{x} \end{array}$$

Fig. 3. Functional operational semantics.

sub term. For example, if $E = f(\bar{v}; \square)$, then $E[(x, y)]$ is $f(\bar{v}; (x, y))$. Together with the `STEP` rule, E determines the evaluation order, where the hole denotes a unique subterm that can be reduced.

A small step reduction $e_1 \longrightarrow e_2$ evaluates e_1 to e_2 . The reduction steps are standard except for always using unboxed tuples to substitute. We write $e[\bar{x}:=\bar{v}]$ for the capture-avoiding substitution of the distinct variables $\bar{x} = (x_1, \dots, x_n)$ with the values $\bar{v} = (v_1, \dots, v_m)$, where n must be equal to m and $\bar{x} \notin \text{fv}(\bar{v})$. When we substitute in a function body, we write $e[\bar{x}:=\bar{v}_1, \bar{y}:=\bar{v}_2]$ to substitute all (distinct) variable \bar{x} and \bar{y} at once, where we again require a capture avoiding substitution with $\bar{x}, \bar{y} \notin (\text{fv}(\bar{v}_1) \cup \text{fv}(\bar{v}_2))$.

When an expression e cannot be reduced further using `STEP`, then either e reduced to an unboxed tuple \bar{v} and we are done, or we call the evaluation *stuck*. We have purposefully described the language and its dynamic semantics without a specific type system, but we can easily define standard Hindley-Milner typing rules [Hindley 1969; Milner 1978] to guarantee that well-typed programs never get stuck.

2.2 FIP: Fully In-Place

As defined, our functional semantics is very liberal and allows expressions that generally cause allocation, like $C x y$. Figure 4 specifies the FIP calculus rules that guarantee that the resulting programs can be evaluated without needing any (de)allocation. The statement $\Delta \mid \Gamma \vdash e$ means that under a borrowed environment Δ and owned environment Γ , the expression e is a well-formed FIP expression. The borrowed environment Δ is a *set* of borrowed variables which generally come from the borrowed parameters of a function f , or by borrowing in the `LET` rule. We write Δ, Δ' for the union of the sets Δ and Δ' . The owned environment Γ is a *multiset* of owned variables, and also *reuse credits*. Following Hofmann [2000a] we denoted these as a diamond \diamond_k , signifying a credit of size k . We can append two owned environments Γ and Γ' as Γ, Γ' . Sometimes, we also write Δ, Γ to join a borrowed set with a multi-set Γ where it is required that there are no reuse credits in Γ , where the result is the borrowed set Δ joined with the elements in Γ . Note that in the current rules, all variables in the Γ environment occur only once as we have no way to duplicate them. In Section 5 we generalize this to the full Perceus calculus.

The FIP rules ensure that variables in the owned environment Γ are used linearly (with some borrowing allowed in `LET`). However, this is a syntactic property and we do not use a linear *type* system. This is much simpler to specify and implement, and also makes FIP independent of any particular type system used by a host language.

The linearity of the FIP calculus is apparent in the `VAR` rule, $\Delta \mid x \vdash x$ where we can only *consume* x if it is the only element of the owned environment Γ . Similarly, the `TUPLE` rule splits the owned environment in n distinct parts, Γ_i , and ensures well-formedness of each constituent value of the tuple, v_i , under the corresponding environment Γ_i . With the `ATOM` rule, $\Delta \mid \emptyset \vdash C$ we can return

FP²: Fully in-Place Functional Programming

$\Gamma ::= \emptyset \mid \Gamma, x \mid \Gamma, \diamond_k$ (owned environment)
 $\Delta ::= \emptyset \mid \Delta, y$ (borrowed environment)

$$\begin{array}{c}
 \frac{}{\Delta \mid x \vdash x} \text{VAR} \qquad \frac{}{\Delta \mid \emptyset \vdash C} \text{ATOM} \\
 \\
 \frac{\Delta \mid \Gamma_i \vdash v_i}{\Delta \mid \Gamma_1, \dots, \Gamma_n \vdash (v_1, \dots, v_n)} \text{TUPLE} \qquad \frac{\Delta \mid \Gamma_i \vdash v_i}{\Delta \mid \Gamma_1, \dots, \Gamma_k, \diamond_k \vdash C^k v_1 \dots v_k} \text{REUSE} \\
 \\
 \frac{\bar{y} \in \Delta, \text{dom}(\Sigma) \quad \Delta \mid \Gamma \vdash e}{\Delta \mid \Gamma \vdash f(\bar{y}; e)} \text{CALL} \qquad \frac{\Delta, \Gamma_2 \mid \Gamma_1 \vdash e_1 \quad \Delta \mid \Gamma_2, \Gamma_3, \bar{x} \vdash e_2 \quad \bar{x} \notin \Delta, \Gamma_2, \Gamma_3}{\Delta \mid \Gamma_1, \Gamma_2, \Gamma_3 \vdash \text{let } \bar{x} = e_1 \text{ in } e_2} \text{LET} \\
 \\
 \frac{y \in \Delta \quad \Delta \mid \Gamma \vdash e}{\Delta \mid \Gamma \vdash y e} \text{BAPP} \qquad \frac{y \in \Delta \quad \Delta, \bar{x}_i \mid \Gamma \vdash e_i \quad \bar{x}_i \notin \Delta, \Gamma}{\Delta \mid \Gamma \vdash \text{match } y \{ C_i \bar{x}_i \mapsto e_i \}} \text{BMATCH} \\
 \\
 \frac{\Delta \mid \Gamma \vdash e}{\Delta \mid \Gamma, \diamond_0 \vdash e} \text{EMPTY} \qquad \frac{\Delta \mid \Gamma, \bar{x}_i, \diamond_k \vdash e_i \quad k = |\bar{x}_i| \quad \bar{x}_i \notin \Delta, \Gamma}{\Delta \mid \Gamma, x \vdash \text{match! } x \{ C_i \bar{x}_i \mapsto e_i \}} \text{DMATCH!} \\
 \\
 \frac{}{\Vdash \emptyset} \text{DEFBASE} \qquad \frac{\Vdash \Sigma' \quad \bar{y} \mid \bar{x} \vdash e}{\Vdash \Sigma', f(\bar{y}; \bar{x}) = e} \text{DEFFUN}
 \end{array}$$

Fig. 4. Well-formed FIP expressions, where the multiplicity of each variable in Γ is 1.

constructors without arguments which we consider allocation free. The owned environment must be empty here since our calculus is not affine: we cannot discard owned variables as that implies freeing a potentially heap allocated value (but in the next section we consider an extension of FIP that allows deallocation as well).

The only way to create a fresh constructor with $k \geq 1$ arguments, is through the `REUSE` rule where we need to check well-formedness of each argument, but we also need a reuse credit \diamond_k to guarantee that the needed space is available at evaluation time. The `DMATCH!` rule creates such reuse credits: we can destructively match on an owned variable x to get a reuse credit \diamond_k in each branch. In each branch the x is no longer owned though (but became a reuse credit instead). For simplicity we only allow matching on a variable in the rules, but we can always rewrite a user expression `match! e { ... }` into `let x = e in match! x { .. }` where x is a fresh owned variable. Again, we do not allow freeing at this point, so reuse credits can only be consumed by the `REUSE` rule or the `EMPTY` rule for zero-sized reuse credits (when an atom is matched).

In contrast, the borrowed match `BMATCH` can only match on borrowed variables and such match can only be used to inspect values without creating fresh reuse credits. Even though variables in the owned environment Γ cannot be discarded (i.e. freed) or duplicated (i.e. shared), we can temporarily *borrow* them. In the `LET` rule the owned environment is split in three parts $\Gamma_1, \Gamma_2, \Gamma_3$. The Γ_1 and Γ_3 environments are passed to e_1 and e_2 respectively, but the Γ_2 environment is passed to e_2 as an owned environment, but also to e_1 as a borrowed environment! Since we *consume* Γ_2 in the derivation of e_2 , we can consider them borrowed in the derivation of e_1 . Note that we still need Γ_3 since Γ_2 is joined with the borrowed Δ environment and as such cannot contain any reuse credits (which can thus be included in Γ_3 instead).

The `CALL` rule is used for a function call $f(\bar{y}; e)$ where we can pass borrowed variables \bar{y} , and the owned argument e . To allow for passing functions, we can also pass a top-level function as part of \bar{y} . In the `BAPP` rule we can call such functions passed as a variable. Since we can only pass them as borrowed, we also only allow borrowed calls of the form $y e$. To prepare for an extension

$$\begin{array}{l}
 S ::= \emptyset \mid S, x \mapsto C^k x_1 \dots x_k \mid S, \diamond_k \\
 E ::= \square \mid C^k x_1 \dots E \dots v_k \mid (x_1, \dots, E, \dots, v_n) \\
 \quad \mid E e \mid x E \mid f(\bar{y}; E) \mid \text{let } \bar{x} = E \text{ in } e \\
 \quad \mid \text{match } E \{ \bar{p} \mapsto e \} \mid \text{match! } E \{ \bar{p} \mapsto e \} \\
 \\
 (let_s) \quad S \mid \text{let } \bar{x} = \bar{z} \text{ in } e \quad \longrightarrow_s \quad S \mid e[\bar{x} := \bar{z}] \\
 (call_s) \quad S \mid f(\bar{y}'; \bar{x}') \quad \longrightarrow_s \quad S \mid e[\bar{y} := \bar{y}', \bar{x} := \bar{x}'] \quad (f(\bar{y}; \bar{x}) = e \in \Sigma) \\
 (anon_s) \quad S \mid (f) \bar{x}' \quad \longrightarrow_s \quad S \mid e[\bar{x} := \bar{x}'] \quad (f(\bar{y}; \bar{x}) = e \in \Sigma) \\
 \\
 (reuse_s) \quad S, \diamond_k \mid C^k x_1 \dots x_k \quad \longrightarrow_s \quad S, x \mapsto C^k x_1 \dots x_k \mid x \quad (k \geq 1, \text{fresh } x) \\
 (atom_s) \quad S \mid C \quad \longrightarrow_s \quad S, x \mapsto C \mid x \quad (\text{fresh } x) \\
 \\
 (bmatch_s) \quad S, y \mapsto C^k \bar{z} \mid \text{match } y \{ \bar{p} \mapsto e \} \quad \longrightarrow_s \quad S, y \mapsto C^k \bar{z} \mid e_i[\bar{y} := \bar{z}] \quad (p_i = C^k \bar{y}) \\
 (dmatch_s) \quad S, x \mapsto C^k \bar{z} \mid \text{match! } x \{ \bar{p} \mapsto e \} \quad \longrightarrow_s \quad S, \diamond_k \mid e_i[\bar{x} := \bar{z}] \quad (p_i = C^k \bar{x})
 \end{array}$$

Fig. 5. Store semantics of FIP.

with full lambda expressions, we only allow owned arguments in a call, as already apparent in the operational semantics. Finally, we can check all top-level functions for well-formedness using the $\Vdash \Sigma$ rule. Any function $f \in \Sigma$ where $\Vdash \Sigma$ is considered *fully in-place*.

Implementing the check algorithmically is straightforward where the owned environment becomes synthesized. For let bindings we first check e_2 and use the synthesized Γ_2, Γ_3 to check e_1 (where Γ_3 only contains reuse credits). When merging synthesized environments we check if linearity is preserved. Since $\Delta \cap \Gamma$ is always empty, we can also infer whether to use a borrowed or destructive match and no such distinction is needed in the user facing syntax – we keep it in our calculus explicit though since we need the distinction in the store semantics.

2.3 Store Semantics

With the FIP calculus defined, we can now define another operational semantics. Figure 5 defines the *store semantics* where we evaluate using a fixed-size store S . The rules in the store semantics all adhere to an important invariant: each step of the evaluation should not allocate or deallocate memory. In this section, we establish a key result relating this store semantics with the functional operational semantics defined previously: under certain conditions, satisfied by all well-formed FIP programs, the store semantics and operational semantics coincide.

The store semantics uses an evaluation context E but this time a full evaluation goes to an unboxed tuple of variables \bar{x} (instead of values \bar{v}). Any constructor is bound in the store S where every element is either a binding $x \mapsto C^k x_1 \dots x_k$ of size k , or a reuse credit \diamond_k of size k .

Using the *eval* rule, we can reduce using small steps in the store semantics. The reduction rules have the form $S \mid e \longrightarrow_s S' \mid e'$ where an e in a store S reduces to e' with a new store S' . The rules are similar to the earlier operational semantics but now we always substitute with variables instead of values. There are now two more rules for evaluating constructor values which are bound in the store. The $(reuse_s)$ transition uses a reuse credit \diamond_k in the store to apply the constructor, while $(atom_s)$ allows atoms to be created freely. The $(bmatch_s)$ and $(dmatch_s)$ reductions differ, where the latter replaces the original constructor binding with a reuse credit of the same size.

Since our store semantics is destructive (in the reuse and dmatch rules), it can fail for expressions where the standard evaluation semantics would succeed. Even for expressions that are well-formed, the store semantics can fail if the initial store has internal sharing. If a shared variable is mutated in place, this would break referential transparency. Thus, we have to require that any variable is

referred to just once in the store – we call this a *linear* store.

Definition 1. (*Store Soundness and Linearity*)

For a store S we write $\text{dom}(S)$ to denote the set of variables x bound in S and write $\text{rng}(S)$ to denote the set of values $C \bar{x}$ bound in S . A store is *sound* if all free variables in $\text{rng}(S)$ are bound: $\text{fv}(\text{rng}(S)) \subseteq \text{dom}(S)$. A store is *linear* if it is sound, and any variable x in $\text{dom}(S)$ occurs at most once in the free variables of $\text{rng}(S)$. By $\text{roots}(S)$ we denote all reuse credits of S and the set of variables in $\text{dom}(S)$ that do not occur in the free variables of $\text{rng}(S)$.

On linear stores mutation is safe; in a reference counted setting such a store corresponds to a heap where all values have a reference count of one. We can now state the main soundness theorem. We write $[S]\bar{x}$ to denote a substitution that recursively replaces variables by their bound value in S . We assume that we are given stores corresponding to the owned and borrowed values, but only require that the store of the owned values is linear. We can then show that the store evaluation leaves the borrowed values unchanged and only modifies the owned values:

Theorem 1. (*The store semantics is sound for well-formed FIP programs*)

If $\Delta \mid \Gamma \vdash e$ and given disjoint stores S_1, S_2 with $\Delta \subseteq \text{dom}(S_1)$, S_1 sound, $\Gamma = \text{roots}(S_2)$ and S_2 linear, then $[S_1, S_2]e \mapsto^* \bar{v}$ implies $S_1, S_2 \mid e \mapsto_s^* S_1, S_3 \mid \bar{x}$ where $[S_3]\bar{x} = \bar{v}$, $\bar{x} = \text{roots}(S_3)$ and S_3 is linear.

This is a strong theorem and the proof is quite involved (see Appendix B), both due to destructive update and the ability to match on variables temporarily borrowed in the `LET` rule. As a corollary, we can now see that any FIP expression can run on the store semantics if we use a store containing the necessary reuse credits, i.e. we give it enough space to allocate upfront:

Corollary 1.

If $e \mapsto^* \bar{v}$ and $\emptyset \mid \diamond_k \vdash e$, then $\diamond_k \mid e \mapsto_s^* S \mid \bar{x}$ and $[S]\bar{x} = \bar{v}$.

Furthermore, any FIP function can be called on stores from the environment:

Corollary 2.

If Σ, f is fully-in-place and $f(\bar{v}_1; \bar{v}_2) \mapsto^* \bar{v}$, and given disjoint linear stores S_1, S_2 with roots \bar{y}, \bar{x} and $[S_1]\bar{y} = \bar{v}_1$, $[S_2]\bar{x} = \bar{v}_2$, then also $S_1, S_2 \mid f(\bar{y}; \bar{x}) \mapsto_s^* S_1, S_3 \mid \bar{z}$, where $[S_3]\bar{z} = \bar{v}$.

We can define the *size* of a store by adding the sizes of all bindings within it. Since atoms and empty reuse credits have size zero, they do not contribute to the size of the store.

Definition 2.

The size $|S|$ of a store S is: $|\emptyset| = 0 \quad |S, \diamond_k| = |S| + k \quad |S, x \mapsto C^k x_1 \dots x_k| = |S| + k$.

With this definition, we can immediately see that the size of the store doesn't change in any reduction of the store semantics. As such, FIP programs can reduce *in-place* without any (de)allocation:

Theorem 2. (*A FIP program reduces in-place.*)

For any $S \mid e \mapsto_s^* S' \mid e'$, we have $|S| = |S'|$.

2.4 FBIP: Allowing Deallocation

Our basic FIP calculus is quite strict and allows neither allocation nor deallocation. We can easily extend it though to allow deallocation. Figure 6 describes the FBIP calculus as an extension of the FIP calculus with deallocation, where the syntax is extended with `drop x`; e to drop an owned variable x , and `free k`; e to free a reuse credit of size k .

The `DROP` rule consumes a variable x from the owned environment. Since the multiplicity of all elements in Γ is still one, this asserts that x is no longer an element of Γ . Similarly, the `FREE` rule allows discarding a reuse credit.

Extended Syntax: $e ::= \dots \mid \text{drop } x; e \mid \text{free } k; e$	Extended evaluation rules: $(dcon_s) \quad S, x \mapsto C^k \bar{x} \mid \text{drop } x; e \longrightarrow_s S \mid \text{drop } \bar{x}; e$ $(free_s) \quad S, \diamond_k \mid \text{free } k; e \longrightarrow_s S \mid e$
$\frac{\Delta \mid \Gamma \vdash e}{\Delta \mid \Gamma, x \vdash \text{drop } x; e} \text{ DROP}$	$\frac{\Delta \mid \Gamma \vdash e \quad k \geq 1}{\Delta \mid \Gamma, \diamond_k \vdash \text{free } k; e} \text{ FREE}$

Fig. 6. The FBIP calculus extends FIP with deallocation.

$\frac{\Delta \mid \Gamma, \diamond_k \vdash e \quad k \geq 1}{\Delta \mid \Gamma, \diamond \vdash e} \text{ INST}$	$\frac{\bar{y}_1 \in \Delta \quad \bar{y}_2 \in \text{dom}(\Sigma) \quad \Delta \mid \Gamma \vdash e \quad f \text{ fip}(n) \quad \bar{y}_2 \text{ fip}(\bar{n})}{\Delta \mid \Gamma, \bar{\diamond}^{n+\bar{n}} \vdash f(\bar{y}_1, \bar{y}_2; e)} \text{ CALL}$
$\frac{\Delta \mid \Gamma \vdash e}{\Delta \mid \Gamma, \diamond \vdash e} \text{ WEAKEN}$	$\frac{\Vdash \Sigma' \quad \bar{y} \mid \bar{x}, \bar{\diamond}^n \vdash e \quad f \text{ fip}(n)}{\Vdash \Sigma', f(\bar{y}; \bar{x}) = e} \text{ DEFFUN}$

Fig. 7. The FIP or FBIP calculus augmented with bounded allocation

The operational semantics is also extended with two new reductions for dropping a bound constructor and freeing a reuse credit. With these new rules and reductions for deallocation, the soundness theorem 1 continues to hold (see Appendix B). Again, we can immediately see that the store semantics now only allows deallocation:

Theorem 3. (A FBIP program can only deallocate.)

For any $S \mid e \mapsto_s^* S' \mid e'$ with the deallocation rules, we have $|S| \geq |S'|$.

In our implementation the `fip` keyword checks if the function is well-formed in the FBIP calculus.

2.5 FIP and FBIP with bounded allocation

We can also model bounded allocation in both the FIP and FBIP calculi to support the `fip`(n) and `fip`(n) keywords, by adding special *allocation* credits \diamond to the Γ environment. These credits do not have a size, and can be used to allocate a cell of any size. Since we usually have several at a time, we write $\bar{\diamond}^n$ for n such credits.

In the `INST` rule, we can allocate a reuse credit \diamond_k of any size from an allocation credit. We do not have to spend allocation credits, which gives them a weakening rule. We modify the `CALL` and `DEFFUN` rules from above to respect the `fip`(n) keyword: When calling a function, we require the necessary allocation credits to call f . Furthermore, we also require allocation credits for all higher-order functions passed to f . Then we can assume in the `BAPP` rule that any allocation credits necessary for the function call have already been paid. Finally, in the `DEFFUN` rule, we obtain as many allocation credits as assigned by the keyword.

It is slightly tricky to infer the necessary allocation credits when also inferring deallocation. For example, there can be several, equally valid inferences for expressions that differ only in their use of `free`:

- $x \mid \diamond \vdash \text{match } x \{ \text{True} \mapsto \text{Cons}(\text{Nil}, \text{Nil}); \text{False} \mapsto \text{Nil} \}$ or
- $x \mid \diamond_2 \vdash \text{match } x \{ \text{True} \mapsto \text{Cons}(\text{Nil}, \text{Nil}); \text{False} \mapsto \text{free } 2; \text{Nil} \}$.

The first expression requires an allocation credit, while the second one requires a reuse credit (which is explicitly deallocated). Which one is preferable depends on the availability of a reuse credit \diamond_2 in the context of e and on the precise annotation on the function containing e : in a `fip`(1) function only the first expression would be valid, while in an `fip` function only the second one

FP²: Fully in-Place Functional Programming

is allowed (if a reuse credit is available). In our implementation, the Γ context is an output of the inference, but we do not create a joint context for all branches of a match immediately. Instead we keep the Γ_i contexts of the branches and resolve them as reuse credits become available from surrounding match-statements.

2.6 Stack Safe FIP

So far, our FIP calculus has allowed us to bound the heap space of the program—but what about the stack space? If we only seek to bound allocations, we could choose to leave it unbounded. In practice, however, the stack space matters: when compiling FIP programs to C we have to assume a relatively small stack and even in a garbage-collected setting growing the stack is not free. To ensure that the stack is bounded, we require two modifications to the calculus. We assume that any function f is defined as part a of mutually recursive group \bar{f} (which might consist of just f or more functions). In the `CALL` rule we then require two additional conditions. Firstly, Σ contains only functions defined before or as part of the current mutually recursive group. Secondly, we constrain the mutually recursive groups \bar{f} by requiring that any recursive calls within this group are tail-recursive. Formally, all function definitions $f \in \bar{f}$ need to be of the form $f(\bar{y}; \bar{x}) = \mathcal{T}[\bar{f}]$:

$$\mathcal{T}[\bar{f}] ::= e_0 \mid f_i(e_0; \bar{e}_0) \mid \text{let } \bar{x} = e_0 \text{ in } \mathcal{T}[\bar{f}] \mid \text{match } e_0 \{ p_i \mapsto \mathcal{T}_i[\bar{f}] \} \mid \text{match! } e_0 \{ p_i \mapsto \mathcal{T}_i[\bar{f}] \} \\ \mid \text{drop } \bar{x}; \mathcal{T}[\bar{f}] \mid \text{free } k; \mathcal{T}[\bar{f}]$$

wher $\bar{f} \not\vdash \text{fv}(e_0)$. In the `CALL` rule, one can pass functions from Σ to the called function. By the first constraint, these functions can only be defined before or mutually recursive with the current definition. In the tail-context we further require that any functions passed as arguments are also not in the mutually recursive group. Thus, we can only pass functions that were defined strictly before the current definition. We call the FIP calculus extended with these requirements FIP^S, in which the stack usage is always bounded. The `fip` keyword in our implementation checks if a function is a well-formed FIP^S function.

In a first-order calculus the previous restrictions immediately imply that the stack size is bounded by the number of mutually recursive groups: any call is either tail-recursive (and executing it does not increase the stack size) or to a previously defined group. However, our FIP calculus is not a first-order calculus since we can pass top-level functions as arguments. Instead, our FIP calculus is a second-order calculus: since we ensure that functions never enter the Γ environment, we cannot return any anonymous function as a result.

In a second-order calculus, it is possible for any top-level function to call back in the list of definitions by calling an anonymous function. However, any sequence of anonymous calls will be to increasingly earlier defined functions. Our stack bound is thus equal to the square of the number of mutually recursive groups.

To show this formally, we use the size of the evaluation context as a proxy for the stack size. We write $|e|$ for the depth of an expression e and $|E|$ for the depth of an evaluation context. We fix a signature Σ and denote by $|e_{\max}|$ the maximum depth of an expression bound in Σ . Then we have:

Theorem 4. (A FIP program uses constant stack space)

Let Σ be fully-in-place such that for all functions f in Σ that are mutually recursive with \bar{f} , we have $f(\bar{y}; \bar{x}) = \mathcal{T}[\bar{f}]$. At any intermediate evaluation step $S \mid f(\bar{y}; \bar{x}) \mapsto_s^* S' \mid E[e']$, we have $|E| \leq |e_{\max}| \cdot |\Sigma|^2$.

This then yields our stack size bound of $|\Sigma|^2$. The additional factor of $|e_{\max}|$ describes the maximum size of the evaluation context within each function. In practice, we would not allocate a stack frame for these parts of the evaluation context. In a first-order context we would expect a stack bound of $|\Sigma|$ (where any function can call functions defined before it). However, in a second-order calculus,

any function can call anonymous functions defined *after* it which adds another factor of $|\Sigma|$ (and see Section Appendix C for a detailed proof).

3 FULLY IN-PLACE TRAVERSALS OVER POLYNOMIAL DATATYPES

A classic example of a fully in-place algorithm is the in-place traversal of a binary tree [Reinking, Xie et al. 2021]. Consider a binary tree with all the values at the tips:

```
type tree<a>
  Bin( left: tree<a>, right: tree<a> )
  Tip( value: a )
```

Similar to our earlier splay tree in Section 1, we can again define a zipper to help traverse the tree in-order:

```
type tzipper<a,b>
  Top
  BinL (up : tzipper<a,b>, right: tree<a>)
  BinR (left : tree<b>, up : tzipper<a,b>)
```

A `tzipper<a,b>` stores fragments of the input tree in-order: those subtrees we have not yet visited are stored using the `BinL` constructor; the subtrees we have already visited are stored in the `BinR` constructor. We can now map over the tree in-order without using heap- or stack space by reusing the `tzipper` nodes. To define the tree map function, we begin by repeatedly stepping down through the input tree to the leftmost tip. Each subtree we have not yet visited, is accumulated in a `BinL` constructor. Once we hit the leftmost leaf, we apply the argument function `f`, and work our way back up, recursively processing any unvisited subtrees:

```
fip fun down( t : tree<a>, ^f : a -> b, ctx : tzipper<a,b> ) : tree<b>
  match! t
    Bin(l,r) -> down( l, f, BinL(ctx,r) ) // go down the left spine, remember to visit r later
    Tip(x)   -> app( Tip(f(x)), f, ctx)   // start upwards along the zipper

fip fun app( t : tree<b>, ^f : a -> b, ctx : tzipper<a,b> ) : tree<b>
  match! ctx
    Top      -> t
    BinR(l,up) -> app( Bin(l,t), f, up) // keep going up rebuilding the tree
    BinL(up,r) -> down( r, f, BinR(t,up) ) // go down a right side

fip fun tmap( t : tree<a>, ^f : a -> b ) : tree<b>
  down(t,f,Top)
```

The mutually tail-recursive `app` and `down` functions are fully in-place since each matched `Bin` can be paired with a `BinL`, each `BinL` with a `BinR`, and finally each `BinR` with a `Bin` again. The definition of `tmap` may seem somewhat involved, yet consider writing this function in an imperative language, without using extra stack- or heap space, mutating pointers throughout the tree.

It is not quite a coincidence that the reuse credits we need in every step line up perfectly – these forms of tail recursive functions can be expressed using a generalisation of zippers known as *dissections* [McBride 2008]. By marking each of these functions as `fip` the compiler checks that they are indeed fully in-place.

Seeing how we can write a map over a binary tree as a FIP function, we may ask if this is possible perhaps for any simple algebraic datatype that can be expressed as a sum of products. It turns out this is indeed the case, and we show this in two steps: first we show in the next subsection a general method for rewriting programs that are tail-recursive *modulo reusable contexts* (TRMReC) such that they are fully in-place. Then, we show how we can generically derive a `map` function for any polynomial inductive datatype to which our TRMReC translation can be applied.

3.1 Tail Recursion Modulo Reusable Defunctionalized CPS Contexts

While our FIP `tmap` function may seem very different from a standard map over trees, it turns out that it actually corresponds to the defunctionalized CPS [Danvy 2008; Reynolds 1972] version of

Translating f to a tail-recursive f' function:

$$\begin{aligned} f(\bar{y}; \bar{x}) &= e \\ f'(\bar{y}; \bar{x}, z) &= \llbracket e \rrbracket_z \\ zipper &= H \mid Z_1 \bar{z}_1 z' \mid \dots \mid Z_n \bar{z}_n z' \end{aligned}$$

Applying the zipper:

$$\begin{aligned} app(\bar{y}; z, \bar{x}') &= match! z \\ &H \rightarrow \bar{x}' \\ &Z_1 \bar{z}_1 z' \rightarrow \llbracket E_1[\bar{x}'] \rrbracket_{z'} \\ &\dots \\ &Z_n \bar{z}_n z' \rightarrow \llbracket E_n[\bar{x}'] \rrbracket_{z'} \end{aligned}$$

Tail recursive contexts with $f \notin \text{fv}(e_0)$:

$$\mathbb{T} ::= \square \mid \text{let } \bar{x} = e_0 \text{ in } \mathbb{T} \mid \text{match } e_0 \{ \overline{p_i \mapsto \mathbb{T}_i} \} \mid \text{match! } e_0 \{ \overline{p_i \mapsto \mathbb{T}_i} \} \mid \text{drop } \bar{x}; \mathbb{T} \mid \text{free } k; \mathbb{T}$$

Tail recursion translation for a function f with zipper z :

$$\begin{aligned} (tctx) \llbracket \mathbb{T}[e] \rrbracket_z &= \mathbb{T}[\llbracket e \rrbracket_z] \\ (base) \llbracket e_0 \rrbracket_z &= app(\bar{y}; z, e_0) \quad \text{where } f \notin \text{fv}(e_0) \\ (tail) \llbracket f(\bar{y}; \bar{x}') \rrbracket_z &= f'(\bar{y}; \bar{x}', z) \\ (ectx) \llbracket E_i[f(\bar{y}; \bar{x}')] \rrbracket_z &= f'(\bar{y}; \bar{x}', Z_i^k \bar{z}_i z) \quad \text{where } \bar{y} \mid \diamond_k, \bar{z}_i \vdash E_i[\square] \text{ and } k = |\bar{z}_i| + 1 \end{aligned}$$

Fig. 8. TRMReC: tail-recursive modulo reusable contexts.

the standard `tmap` function:

```
fun tmap(t : tree<a>, ^f : a -> b) : tree<a>
  match! t
    Bin(l,r) -> val l' = tmap(l,f) in val r' = tmap(r,f) in Bin(l', r')
    Tip(a)   -> Tip(f(a))
```

Let us focus on the first branch, where a CPS-translation yields the following closures:

```
Bin(l,r) -> tmap(l,f, fn(l'){ tmap(r,f, fn(r'){ k(Bin(l', r')) }) })
```

Comparing with our `tzipper` type, we can identify `Top` with the identity function, `BinR` with the inner closure (`fn(r') k(Bin(l', r'))`), and `BinL` with the outer closure – the zipper is just the defunctionalization of the closures:

```
fn(x) x          === Top          -> t
fn(r') k(Bin(l',r')) === BinR(l',k) -> app(Bin(l',t), f, k)
fn(l') tmap(r,f,fn(r'){ k(Bin(l',r')) }) === BinL(k,r) -> down(r,f,BinR(t,k))
```

The arguments r' and l' to the closures correspond to the tree t , the `down` function to the transformed `tmap` function, and `app` applies the defunctionalized continuation k to the new tree. As shown by Danvy [2022], this defunctionalized CPS-transformation applies widely and can transform many programs from direct-style to tail-recursive form. But are these techniques also applicable when writing FIP programs? Sobel and Friedman [1998] show that it is always possible to reuse the zipper for the result in all anamorphisms. In fact, in the above translation, it is even possible to reuse the initial tree to construct the zipper.

Using this insight, we can give a general translation to tail-recursive programs that is guaranteed fully in-place. It is inspired by the defunctionalized CPS-translation, but we have to make several small adjustments to make it work. For example, notice how the borrowed function f is not included in the zipper, but instead passed directly to `app`. This is crucial, since we can not store a borrowed value inside a data structure. Figure 8 shows the formal transformation, based on the general framework of tail recursion *modulo context* as shown recently by Leijen and Lorenzen [2023].

Starting with a function $f(\bar{y}; \bar{x}) = e$, we first define the zipper by creating a constructor `H` for the identity and one constructor Z_i each for each evaluation context E_i in e that contains a recursive call to f . Each constructor carries the free variables \bar{z}_i of its evaluation context and the link to the parent zipper z' . We transform f into $f'(\bar{y}; \bar{x}, z)$ and provide an $app(\bar{y}; z, \bar{x}')$ function, where we ensure that both receive the same borrowed variables \bar{y} . The calls to f' receive the current zipper

as an extra argument z and is defined by the translation $\llbracket e \rrbracket_z$ below. The *app* function matches on the current zipper (as before) and resumes execution in the relevant (transformed) evaluation context.

The transformation follows defunctionalized CPS contexts of the TRMC framework [Leijen and Lorenzen 2023, Sec. 4.3] with the (*tctx*), (*base*), (*tail*), and (*ectx*) rules. If we encounter a tail context \mathbb{T} , we continue the transformation in every hole using the (*tctx*) rule. When we encounter a term e_0 which has no recursive calls, the (*base*) rule inserts a call to *app* to apply the result of e_0 to the continuation stored in z . If we encounter a tail-call we simply leave it as is. Finally, if we find a call in an evaluation context, we turn it into a tail-call by storing the free variables of E_i and the current zipper. Notice that we cannot have a tail-context nested in an evaluation context as the translation assumes that programs are in A-normal form [Flanagan et al. 1993].

So far this transformation describes just how to enable tail recursion on general defunctionalized CPS contexts but the result is not yet guaranteed to be FIP. To guarantee that reuse applies we need to preserve the side-condition on (*ectx*), which ensures that E_i does not depend on borrowed variables other than \bar{y} (which could not be stored in an accumulator) and that there is a space credit of the appropriate size for z and the free variables \bar{z}_i of E_i .

If this condition is met, this transformation yields a tail-recursive, fully in-place program:

Theorem 5. (*The TRMReC transformation is sound.*)

Let f be a function with $\bar{y} \mid \bar{x} \vdash f(\bar{y}; \bar{x})$ and $f(\bar{v}_1; \bar{v}_2) \longrightarrow^* \bar{w}$. If it can be transformed into f' , then $\bar{y} \mid \bar{x}, z \vdash f'(\bar{y}; \bar{x}, z)$ and $\bar{y} \mid z, \bar{x} \vdash \text{app}(\bar{y}; z, \bar{x})$ and $f'(\bar{v}_1; \bar{v}_2, H) \longrightarrow^* \bar{w}$.

See Appendix D for the proof. We also generalize this theorem to handle recursive calls with varying borrowed arguments and the case where $k \geq |\bar{z}_i| + 1$ (as common in folds) or more than one reuse credit is available. Clearly, this translation can apply to the tmap introduced at the start of this section. We just have to check the side condition:

- For $E_1 := \text{let } l' = \square \text{ in let } r' = \text{tmap}(f; r) \text{ in Bin } l' r'$ we have $f \mid \diamond_2, r \vdash E_1[\square]$.
- For $E_2 := \text{let } r' = \square \text{ in Bin } l' r'$ we have $f \mid \diamond_2, l' \vdash E_2[\square]$.

Thus the condition is fulfilled and the translation succeeds.

An important special case of the more general transformation is the tail recursion modulo (product) cons transformation [Bour et al. 2021; Leijen and Lorenzen 2023], where $E = \text{let } \bar{x} = \square \text{ in } \bar{v}$. This can be compiled more efficiently by allocating the values \bar{v} upfront and storing E as pointers to the slots \bar{x} in \bar{v} . In our programming examples in section 4 we assume this optimization, but in systems that lack it one can instead obtain a FIP program by using reusable contexts.

3.2 Schorr-Waite Tree Traversals

Using the translation in the previous section we can now generalize the tmap function to any polynomial inductive datatype. Following the approach by van Laarhoven [2007] to generically derive functors in Haskell, we use a generic macro $\$map_\tau$ to define a (non tail-recursive) *map* function for any type $T \alpha$ in a straightforward way, where we match on each constructor in T and call the $\$map$ macro on each of the fields:

$$\begin{array}{ll}
 \text{map}(f; x) = \text{match! } x & \$map_{T \alpha}(f; x_i) = \text{map}(f; x_i) \\
 C^k(x_1 : \tau_1) \dots (x_k : \tau_k) \rightarrow & \$map_\alpha(f; x_i) = f(x_i) \\
 \text{let } y_1 = \$map_{\tau_1}(f; x_1) \text{ in} & \$map_\tau(f; x_i) = x_i \quad \text{otherwise} \\
 \dots & \\
 \text{let } y_k = \$map_{\tau_k}(f; x_k) \text{ in} & \\
 C^k y_1 \dots y_k & E_i = \text{let } y_i = \square \text{ in } \dots \text{ in } C^k y_1 \dots y_k \\
 \dots & f \mid \diamond_k, y_1, \dots, y_{i-1}, x_{i+1}, \dots, x_k \vdash E_i[\square]
 \end{array}$$

FP²: Fully in-Place Functional Programming

The $\$map_\tau$ macro dispatches on the *type* of its argument. if x has type $T \alpha$ it generates a recursive call $map(f; x)$ if x has type α , it generates a call $f(x)$; otherwise it leaves the argument unchanged. In this definition, all recursive calls to *map* happen in $\$map$, each application of which is in an evaluation context E_i where the side-condition holds. Thus, we can apply the TRMReC transformation of the previous section and automatically obtain a tail-recursive fully in-place version of *map*.

Why does reuse work so naturally here? Part of the solution seems to be that the link to the parent is stored together with the other free variables. In contrast, McBride [2008] defines a generic *fold* function which is not fully in-place since it stores the defunctionalized continuations on a stack. Nevertheless, as McBride [2001] shows, the defunctionalized continuations correspond to the (generalised) *derivative* of a regular type $T \alpha$. For every constructor C with k recursive subtrees, the derivative datatype has k constructors, one for each possible continuation. Reuse then arises naturally, as the constructors of the derivative and original datatype can line up perfectly.

In the literature on imperative algorithms, these traversals that use no extra stack space except for direction hints (as encoded in the constructors of the zipper datatype), are known as Schorr-Waite traversals [Schorr and Waite 1967]. Effectively, we can thus derive a Schorr-Waite traversal for any polynomial algebraic datatype and use the reuse analysis of the FIP calculus to compile it to the corresponding imperative code. This is remarkable as imperative Schorr-Waite traversals are notoriously difficult to get right or prove correct. In his famous work on separation logic, Reynolds [2002] writes:

The most ambitious application of separation logic has been Yang’s proof of the Schorr-Waite algorithm for marking structures that contain sharing and cycles.

Of course, our construction cannot handle cycles, but rather shows the traversal of trees (or any polynomial inductive datatype in general). The tree traversal by itself, however, is already quite complicated and has become a benchmark for verification frameworks [Loginov et al. 2006; Walker and Morrisett 2000]. Furthermore, our translation also shows that the Schorr-Waite tree traversal is equivalent to a stack-based depth-first traversal (like the standard $tmap$). This was already shown by Yang [2007] in the context of separation logic, but that required more advanced methods than straightforward induction.

4 FURTHER EXAMPLES OF FULLY IN-PLACE ALGORITHMS

Many common functions used in functional programming are already FIP in their standard definition, like *map* or *reverse*. In this section we want to present some advanced examples that test the limits of what is possible. Along the way, we see several techniques that may be of general use for designing algorithms in a language with in-place reuse. These include passing reuse credits to functions, padding constructors and the *partition* datastructure. Full listings of the examples in this section can be found in Appendix A.

4.1 Imperative red-black tree insertion, functionally

Can insertion into a red-black tree be FIP? The traditional implementation of red-black trees, due to Okasaki [1999], can indeed be written to use all its arguments in-place. However, it occasionally has to rebalance the result of the recursive call and thus uses stack space linear in the depth of the tree. However, we can avoid this by using a zipper (as in Section 3) and balancing while reconstructing the tree from the zipper. Surprisingly, this yields a functional implementation which is almost identical to the imperative red-black tree insertion algorithm described in the popular “Introduction to Algorithms” textbook [Cormen et al. 2022]. We first define the tree and its zipper:

```

type color { Red; Black }
type tree<k,v>
  Node(c : color, l : tree<k,v>, k : k, v : v, r : tree<k,v>)
  Leaf
type accum<k,v>
  Done
  NodeL(c : color, l : accum<k,v>, k : k, v : v, r : tree<k,v>)
  NodeR(c : color, l : tree<k,v>, k : k, v : v, r : accum<k,v>)

```

This definition is typical for a functional datastructure, but different from the imperative definition as a node does not store a pointer to its parent. Such parent-pointers would not work well in our setting, because they create many small cycles between parent and children, which can never be in-place reused again. Instead, we use pointer-reversal along the insertion path. On this path, the zipper contains the parent links *instead of* children links and no cycles are created.

We can insert a value into the tree by recursing into the left- or right-subtree until we either find the key with an existing value or a leaf. In the latter case, we will have to allocate a new node for key and value. During the recursion, we turn the nodes of the tree into the `accum` zipper. At the end, we thus have a subtree with our new value and a zipper. We create a rebalanced red-black tree by calling the `fixup` function below on them.

But how do we write `fixup`? Thankfully, we can translate it almost verbatim from the `rb-insert-fixup` procedure in section 13.3 of Cormen et al. [2022]! We present the code here to illustrate that this function is FIP and closely follows its imperative counterpart. A detailed explanation can be found in [Cormen et al. 2022]. The function distinguishes three cases (marked on the left of the function definition) that correspond to the three cases in the textbook implementation. Case one translates directly (even if we have it twice). The second case is the most complicated, rotating an inner part left, before rotating the outer part right. In case three, it is possible to stop `fixup` (by calling `rebuild` defined below) as the parent node was colored black.

```

fix fun fixup( zp : accum<k,v>, z : tree<k,v> ) : tree<k,v>
  match! zp
  NodeL( Red, zpp, zpk, zpv, zpr ) -> match zpp
    NodeL( Black, zppp, zppk, zppv, y ) ->
      if is-red(y)
      then fixup(zppp, Node( Red, Node( Black, z, zpk, zpv, zpr ), zppk, zppv, y.set-black)) // (1)
      else rebuild(zppp, right-rotate( Node( Red, Node( Black, z, zpk, zpv, zpr ), zppk, zppv, y))) // (3)
    NodeR( Red, zpl, zpk, zpv, zpp ) -> match zpp
      NodeL( Black, zppp, zppk, zppv, y ) ->
        if is-red(y)
        then fixup(zppp, Node( Red, Node( Black, zpl, zpk, zpv, z ), zppk, zppv, y.set-black)) // (1)
        else rebuild(zppp, right-rotate(
          Node( Red, left-rotate( Node( Red, zpl, zpk, zpv, z.set-black ), zppk, zppv, y))) // (2)
        // and cases as above with "left" and "right" interchanged
        _ -> rebuild(zp, z)

```

We have left some functions unspecified: `set-black` sets the color of a node to black. `is-red` returns a boolean indicating whether the given node is red. We have to make its argument borrowed so that we apply `r2` to it. The left rotation is as usual (and the right rotation its mirror image).

```

fix fun is-red(^t : tree) : bool
  match t { Node( Red ) -> True; _ -> False }

fix fun left-rotate(t : tree<k,v>) : tree<k,v>
  match! t { Node(c,l,k,v, Node(c1,l1,k1,v1,r1)) -> Node(c1, Node(c,l,k,v,l1), k1, v1, r1); t' -> t' }

```

The last remaining function is `rebuild`, which is called when balancing is finished. In the imperative implementation, this simply marks the root black and returns the root. But in our version, the root is now hidden in the zipper and we have to rebuild the tree from the zipper (without balancing) to access the root:

FP²: Fully in-Place Functional Programming

```
fip fun rebuild(z : accum<k,v>, t : tree<k,v>)
  match! z
  NodeR(c, l, k, v, z1) -> rebuild(z1, Node(c, l, k, v, t))
  NodeL(c, z1, k, v, r) -> rebuild(z1, Node(c, t, k, v, r))
  Done -> t.set-black
```

In practice, the imperative version benefits from not having to rebuild the tree (and fixup considers some cases specifically to be able to enter rebuild earlier). Thus we can not quite achieve the same efficiency in a functional version. However, our version can also be used persistently (see Section 5) and might be easier to understand.

4.2 Sorting lists in-place

Is it possible to run merge sort in-place? The traditional functional implementation first turns each element into a singleton list. A singleton list is obviously sorted, so we can now pairwise merge such sorted lists. Finally, we end up with just one sorted list, which we extract:

```
[4,3,2,1] -> [[4],[3],[2],[1]] // create sorted singleton lists
-> [[3,4],[1,2]] -> [[1,2,3,4]] // merge pairs of sorted lists
-> [1,2,3,4] // extract sorted list
```

Here, the output takes up just as much space as the input - so a fully in-place implementation might be possible. But in the first step, many singleton lists are created for which no reuse tokens are available, so the traditional implementation is *not* fully in-place.

However, we can make it FIP by using a tailor made datastructure, that can store the sorted sublists while taking up exactly the same amount of space as the original list. Our partition is a list partitioned into sublists, which are either *Ones* or *Subs* of at least two elements. We exploit that we have at least two elements by storing two elements in the last cell of a *list2*.

```
type partition<a>
  Sub(list : list2<a>, tail : partition<a>)
  One(elem : a, tail : partition<a>)
  End
type list2<a>
  Cons2(x : a, tail : list2<a>)
  Nil2(x : a, y : a)
```

Considering the memory usage, we see that a *list2* can store n elements in $n - 1$ cells. As a result, a partition of n elements uses just as much space as a list of n elements, while also keeping information about the partitioning. With that in hand, we can implement an in-place list mergesort:

```
Cons(4,Cons(3,Cons(2,Cons(1,Nil)))) // start with unsorted list
-> One(4,One(3,One(2,One(1,End)))) // created sorted singleton lists
-> Sub(Nil2(3,4),Sub(Nil2(1,2),End)) // merge pairs of sorted lists
-> Sub(Cons2(1,Cons2(2,Nil2(3,4))),End) // ... until only one list2 is left
-> Cons(1,Cons(2,Cons(3,Cons(4,Nil)))) // convert back to list
```

Note that this does not solve the (much harder) in-place sorting problem for arrays, where we can not distinguish sublists through tags and shuffling pointers.

This datastructure is also helpful to implement a quicksort that can not run out of stack. The typical functional but in-place implementation is [Baker 1994a; Hofmann 2000b; Hudak 1986]:

```
fbip fun quicksort(xs : list<a>)
  match! xs
  Nil -> Nil
  Cons(pivot, xx) ->
    val (lo, hi) = split(pivot, xx) // split xx into (lo,hi) in-place
    val (lo', hi') = (quicksort(lo), quicksort(hi)) // sort the sublists
    append(lo', Cons(pivot, hi'))
```

This code is not FIP because the stack usage is not bounded and might grow linear with the length of the input list. Of course, we could apply the defunctionalized CPS-transformation (see Section 3), but our side-condition fails:

```

Cons(pivot, xx) -> // reuse credit of size 2 available
  val (lo, hi) = split(pivot, xx)
  quicksort'(lo, Z1(pivot, hi, zipper)) // reuse credit of size 3 needed

```

The problem here is that the zipper needs to store `pivot,hi` and the parent zipper, which requires more space than we have available. This is because, in a stack-safe quicksort, the zipper needs to keep track of all the pivots and `hi` lists that still need to be sorted. However, we can use a partition structure as the zipper where we store the pivots as singletons and the `hi` either not at all (if `hi` is empty), as a singleton (if `hi` is a one-element list) or else as a `list2`. We can pass the parent zipper into the `split` function, which now returns a list `lo` and a partition `hi` which includes the zipper. Then we obtain a fully in-place solution:

```

Cons(pivot, xx) -> // reuse credit of size 2 available
  val (lo, hi) = split(pivot, xx, zipper)
  quicksort'(lo, One(pivot, hi)) // reuse credit of size 2 needed

```

4.3 Finger trees

Finally, as an advanced example, we want to consider finger trees [Claessen 2020; Hinze and Paterson 2006], an efficient functional implementation of sequences. Yet, at first glance the `cons` function on finger trees does not appear to be FIP: only the `More` constructors can be reused as the other datatypes do not match up for reuse. We can fix this, however, by *padding* all constructors with a dummy atom `Pad` so that they all have three slots.

We use Claessen [2020]’s version of finger trees, since Hinze and Paterson’s version has an extra `Four` constructor which would require us to pad *all* constructors to four fields.

```

fun cons(x : a, s : seq<a>) : seq<a>
  match! s
  Empty -> Unit(x, Pad, Pad)
  Unit(y, _, _) -> More(One(x, Pad, Pad), Empty, One(y, Pad, Pad))
  More(One(y, _, _), q, u) -> More(Two(x, y, Pad), q, u)
  More(Two(y, z, _), q, u) -> More(Three(x, y, z), q, u)
  More(Three(y, z, w), q, u) -> More(Two(x, y, Pad), cons(Pair(z, w, Pad), q), u)

```

We have gotten rid of all deallocations since all constructors on the left of `->` can be paired with one to the right. But we still have allocations in the `Empty`, `Unit` and `More(Three)` cases. Even worse, the `cons` can recurse up to $O(\log n)$ -times in the `More(Three)` case and require a new memory cell each time, so this function is not `fip(n)` or `fbip(n)`. However, this case is very unlikely as the amortized complexity analysis of finger trees shows that `cons` only recurses $O(1)$ -times on average and thus only uses a constant amount of memory.

Therefore, we pair a finger tree `seq<a>` with a buffer which contains exactly the memory needed. Our buffer is just a padded list of size 3, which makes it available for reuse with the rest of the finger tree.

```

type buffer { BEmpty; BCons(next : buffer, b : pad, c : pad) }

```

We then pass the necessary reuse credits of size 3 to `cons`, which we either use to create a new cell in the finger tree or fill up the buffer. If we recurse into `cons`, we draw the necessary memory back from the buffer. Then we only need to ensure that we pass enough credits so that the buffer is never empty. Inserting two elements `x,y` into an empty finger tree yields `More(One(x, Pad, Pad), Empty, One(y, Pad, Pad))`, so it would seem that we need to pass at least two credits. But that would mean that we need $6n$ space to represent n elements in a finger tree! We can do better by specializing `More(One)` as `More0` to represent the two-element list as `More0(x, Empty, One(y, Pad, Pad))`. With this modification, it suffices to pass in a single reuse credit per element for a space overhead of $3n$ space, which is close to the $2n$ factor of singly-linked lists. Our `cons` function then takes a reuse credit `unit3` and becomes:

Extended syntax :

$$v ::= \dots \mid \lambda \bar{x}. e \quad (\text{lambda with } \bar{z} = \text{fv}(\lambda \bar{x}. e))$$

$$e ::= \dots \mid \text{dup } x; e \mid \text{dropru } x; e \mid \text{alloc } k; e$$

Extended evaluation steps :

$$(\text{beta}) \quad (\lambda \bar{x}. e) \bar{v} \longrightarrow e[\bar{x} := \bar{v}]$$

$$(\text{dup}) \quad \text{dup } x; e \longrightarrow e$$

$$(\text{dropru}) \quad \text{dropru } x; e \longrightarrow e$$

$$(\text{alloc}) \quad \text{alloc } k; e \longrightarrow e$$

$$\frac{\Delta \mid \Gamma, x \vdash e \quad x \in (\Delta, \Gamma)}{\Delta \mid \Gamma \vdash \text{dup } x; e} \text{ DUP}$$

$$\frac{\Delta \mid \Gamma, \diamond_k \vdash e \quad k = \text{size}(x)}{\Delta \mid \Gamma, x \vdash \text{dropru } x; e} \text{ DROPRU}$$

$$\frac{\Delta \mid \Gamma, \diamond_k \vdash e \quad k \geq 1}{\Delta \mid \Gamma \vdash \text{alloc } k; e} \text{ ALLOC}$$

$$\frac{x \in (\Delta, \Gamma) \quad \Delta \mid \Gamma, \bar{x}_i \vdash e_i}{\Delta \mid \Gamma \vdash \text{match } x \{ C_i \bar{x}_i \mapsto \text{dup } \bar{x}_i; e_i \}} \text{ MATCH}$$

$$\frac{\Delta \mid \Gamma_1 \vdash e_1 \quad \Delta \mid \Gamma_2 \vdash e_2}{\Delta \mid \Gamma_1, \Gamma_2 \vdash e_1 e_2} \text{ APP}$$

$$\frac{\emptyset \mid \bar{z}, \bar{x} \vdash e \quad \bar{z} = \text{fv}(\lambda \bar{x}. e)}{\Delta \mid \bar{z} \vdash \lambda \bar{x}. e} \text{ LAM}$$

Fig. 9. The λ^{fip} calculus extends the Perceus linear resource calculus with borrowing reuse, and unboxed tuples. The calculus extends the syntax, rules, and functional semantics of the FBIP calculus as shown in Figure 6 and 4. The multiplicity of each variable in Γ is unconstrained.

```

fip fun cons(x : a, u3 : unit3, s : seq<a>, b : buffer) : (seq<a>, buffer)
  match! s
    Empty -> (Unit(x, Pad, Pad), b)
    Unit(y, _, _) -> (More0(x, Empty, One(y, Pad, Pad)), b)
    More0(y, q, u) -> (More(Pair(x, y, Pad), q, u), b)
    More(Pair(y, z, _) , q, u) -> (More(Triple(x, y, z), q, u), BCons(b, Pad, Pad))
    More(Triple(y, z, w), q, u) ->
      match! b
        BCons(b', _, _) ->
          val (q', b'') = cons(Pair(z, w, Pad), u3, q, b')
          (More(Pair(x, y, Pad), q', u), b'')

```

This function is now fully in-place. In the `More(Two)` case we store an unneeded credit in the buffer. In the `More(Three)` case we recurse⁴ and take a reuse credit from the buffer. The buffer has the invariant that, given n_1 `Triple`, n_2 `Three` and n_3 `Two` constructors in the finger tree, its size is just $n_1 + 2 * n_2 + n_3$. Since this invariant is maintained in the `cons` function, the buffer is never empty.

5 REFERENCE COUNTING WITH BORROWING AND UNBOXING

In this section we formalize the connection between the FIP calculus and Perceus precise reference counting [Lorenzen and Leijen 2022; Reinking, Xie et al. 2021]. Our implementation of FIP in Koka uses this approach where detection of the uniqueness of owned arguments happens dynamically at run-time. Figure 9 formalizes the λ^{fip} calculus as an extension of the syntax and operational semantics of the FBIP calculus (in Figure 4 and 6). It has full lambda expressions $\lambda \bar{x}. e$ now since arbitrary allocation is allowed. Here we write the free variables of the lambda expression explicitly as (the multiset) \bar{z} . This is not needed for the functional operational semantics but as we see later, we require it for the heap based operational semantics. Moreover, we have a `dup x; e` and `dropru x; e` expressions that let us duplicate owned variables and explicitly reuse dropped variables. Finally, the `alloc k; e` allows for arbitrary allocation of constructors by creating reuse credits \diamond_k at runtime.

The *(beta)* evaluation rule for lambda expressions is standard, and we can see that the *(dup)*, *(dropru)*, and *(alloc)* rules have no effect in the functional operational semantics (and are only used

⁴The recursive call is here in tail-position modulo product-contexts [Leijen and Lorenzen 2023], which can be efficiently compiled to a tail-recursive call. However, we could also apply the TRMReC transformation to eliminate it (see section Appendix D).

in the heap semantics).

We rephrase the original Perceus linear resource calculus, called λ^1 [Reinking, Xie et al. 2021], in Figure 9 as a *type system* (instead of a typed translation). We call any expressions in $\Delta \mid \Gamma \vdash e$ well-formed, and such expression always uses correct reference counting when evaluated, i.e. it never drops a value from the heap that is still needed later, or leaves garbage in the heap at the end of an evaluation. Moreover, the new type rules also extend the original rules with borrowing and unboxed tuples, and give a characterization of reuse based on reuse credits.

Just like the FIP calculus, the rules are still based on linear logic with a linear owned Γ environment, but unlike a pure linear logic it now has an escape hatch: through rules like `DUP` we can freely duplicate “linear” variables by maintaining reference counts dynamically at runtime. As we can see, there is a suprisingly close connection between λ^{fip} and the FBIP and FIP calculi where each one is a strict subset of the other: $\text{FIP} \subset \text{FBIP} \subset \lambda^{\text{fip}}$. As such, the FIP calculus is exactly the subset of λ^{fip} that excludes the rules that require dynamic reference counting!

The `DUP` rule is the rule that either allows us to use a borrowed variable ($x \in \Delta$) as owned, or duplicates an owned variable ($x \in \Gamma$). The `ALLOC` rule now allows arbitrary allocation of a constructor by adding a reuse credit \diamond_k to the owned environment. With the full lambda expressions, we also have an `APP` rule to apply an argument to a lambda expression. Here, we split the owned environment in two parts for each subexpression. We could have been more elaborate and allow borrowing of Γ_2 in the e_1 derivation just like our earlier `LET` rule in Figure 4. We refrain from doing that here for simplicity as we can always use `LET` if borrowing is required.

The `LAM` rule requires that all free variables of the lambda expression are owned (which are needed to create the initial closure). In the body, we check with the free variables (from the closure) and the passed in parameters as all owned. Borrow information is not part of a type, so only top-level functions can take borrowed arguments (using the `CALL` rule).

The `MATCH` rule can match on any borrowed or owned variable. However, each branch must start by dupping the matched constructor fields (as `dup(\bar{x}_i)`). Indeed, since the match is non-destructive each field is now reachable directly but also via the original x (and thus we need to increment the reference count at runtime). For simplicity, the `MATCH` rule can only match variables but we can always rewrite an expression `match e { ... }` into `let x = e in match x { ... }` for a fresh x when required. Since `MATCH` no longer creates reuse credits, we can now create them explicitly instead using the “drop reuse” `DROPRU` rule. This drops a variable x , and immediately allows for a reuse credit \diamond_k where k is the allocated size of x .

With the new `MATCH` and `DROPRU` rules we no longer require the destructive match and corresponding rule of the FIP calculus and we can always replace any destructive match:

`match! x { Ci \bar{x}_i → ei }` with `match x { Ci \bar{x}_i → dup \bar{x}_i; dropru x; ei } .`

In particular, if the FIP `match!` expression is well-formed, we have:

$$\frac{\Delta \mid \Gamma, \bar{x}_i, \diamond_k \vdash e_i \quad (1)}{\Delta \mid \Gamma, x \vdash \text{match! } x \{ C_i \bar{x}_i \rightarrow e_i \}} \text{DMATCH!}$$

and thus we can also derive that the translated match is well-formed in the λ^{fip} calculus:

$$\frac{x \in \Gamma, x \quad \frac{\Delta \mid \Gamma, \bar{x}_i, \diamond_k \vdash e_i \quad (1)}{\Delta \mid \Gamma, x, \bar{x}_i \vdash \text{dropru } x; e_i} \text{DROPRU}}{\Delta \mid \Gamma, x \vdash \text{match } x \{ C_i \bar{x}_i \rightarrow \text{dup } \bar{x}_i; \text{dropru } x; e_i \}} \text{MATCH}$$

Furthermore, unlike the FIP or FBIP calculus, we can always elaborate a plain expression with `dup`, `drop`, `free`, `alloc`, and `dropru` to make it a well-formed λ^{fip} expression. The heap semantics can thus always be used to evaluate an expression. In particular, we can easily adapt the Perceus algorithm [Reinking, Xie et al. 2021] to elaborate plain expressions with correct reference count

FP²: Fully in-Place Functional Programming

A heap H extends a store S with reference counts $n \geq 1$, and can hold closures as well:

$$H ::= \emptyset \mid H, \diamond_k \mid H, x \mapsto^n \varphi \quad \text{where } \varphi ::= C \bar{x} \mid \lambda \bar{x}. e$$

The \mapsto_h relation extends the \mapsto_s (using H for S) relation:

$$\begin{array}{l}
(\text{dup}_h) \quad H, x \mapsto^n v \quad \mid \text{dup } x; e \quad \longrightarrow_h \quad H, x \mapsto^{n+1} v \quad \mid e \\
(\text{dlam}_h) \quad H, x \mapsto^1 \lambda \bar{x}. e' \quad \mid \text{drop } x; e \quad \longrightarrow_h \quad H \quad \mid \text{drop } \bar{z}; e \\
(\text{drop}_h) \quad H, x \mapsto^{n+1} v \quad \mid \text{drop } x; e \quad \longrightarrow_h \quad H, x \mapsto^n v \quad \mid e \quad (\text{if } n \geq 1) \\
(\text{dropru}_h) \quad H, x \mapsto^{n+1} C^k \bar{x} \quad \mid \text{dropru } x; e \quad \longrightarrow_h \quad H, \diamond_k, x \mapsto^n C^k \bar{x} \quad \mid e \quad (\text{if } n \geq 1) \\
(\text{dconru}_h) \quad H, x \mapsto^1 C^k \bar{x} \quad \mid \text{dropru } x; e \quad \longrightarrow_h \quad H, \diamond_k \quad \mid \text{drop } \bar{x}; e \\
(\text{alloc}_h) \quad H \quad \mid \text{alloc } k; e \quad \longrightarrow_h \quad H, \diamond_k \quad \mid e \\
(\text{lam}_h) \quad H \quad \mid \lambda \bar{x}. e \quad \longrightarrow_h \quad H, x \mapsto^1 \lambda \bar{x}. e \quad \mid x \quad (\text{fresh } x) \\
(\text{app}_h) \quad H \quad \mid (y) \bar{y} \quad \longrightarrow_h \quad H \mid \text{dup } \bar{z}; \text{drop } y; e[\bar{x} := \bar{y}] \quad (y \mapsto^n \lambda \bar{x}. e \in H)
\end{array}$$

Fig. 10. Heap semantics of λ^{fip} – extending the FBIP store semantics as shown in Figure 5 and 6.

instructions.

5.1 Elaborating borrowing

However, the elaboration needs to take some care when functions with borrowed parameters are passed as arguments. For example, what should happen when the user passes `is-red` to `map`:

```
map(is-red, [t1, t2, t3]) : list<bool>

fip fun is-red(^t : tree) : bool
fip fun map( l : list<a>, ^f : a -> b ) : list<b>
```

Here, `map` assumes that the given function f takes its argument as owned, while `is-red` takes its argument t as borrowed. In a linear type system (and the FIP calculus), such a program would be rejected. However, we follow a different approach here: Since we seek to integrate fully in-place programs seamlessly into a general-purpose functional programming language, the program above should be accepted. Instead, we *wrap* functions with borrowed parameters, into a closure that takes only owned parameters, calls the original function and then drops as necessary [Ullrich and de Moura 2019]. Formally, we can derive:

$$\frac{\frac{\bar{y} \mid \bar{x} \vdash \bar{x}}{\bar{y} \mid \bar{x} \vdash f(\bar{y}; \bar{x})} \text{CALL} \quad \frac{\emptyset \mid z \vdash z}{\emptyset \mid \bar{y}, z \vdash \text{drop } \bar{y}; z} \text{DROP}}{\emptyset \mid \bar{y}, \bar{x} \vdash \text{let } z = f(\bar{y}; \bar{x}) \text{ in drop } \bar{y}; z} \text{LET}}{\Delta \mid \emptyset \vdash \lambda^{\emptyset} (\bar{y}, \bar{x}). \text{let } z = f(\bar{y}; \bar{x}) \text{ in drop } \bar{y}; z} \text{LAM}$$

This wrapping can be done cheaply for top-level functions with borrowed parameters, since the new closure has no free variables and can be lambda-lifted to the top-level without allocation. Later compilation stages may then decide to inline the nested function call – making the wrapping zero-overhead at the cost of code duplication.

However, this would be more expensive on closures, where f itself is a free variable! Then we would have to allocate the wrapping closure and could not rely on inlining to remove the overhead of the extra indirection. For this reason, borrowing is usually not done on closures – and we reflect this in our FIP and λ^{fip} calculi to remain compatible with this interpretation of reference counting.

5.2 Heap semantics

Figure 10 gives a heap based operational semantics for our Perceus calculus. Here we generalize the store S from the FIP calculus (Figure 5) to contain a reference count $n \geq 1$ for each binding. The heap now contains reuse credits \diamond_k , a constructor binding $x \mapsto^n C \bar{x}$, or a closure $x \mapsto^n \lambda \bar{x}. e$.

We extend the original FBIP store semantics in Figure 5 and Figure 6 with new rules where the evaluation context and eval rule stays the same (just replacing a store S with a heap H). For the (bmatch) rule we can allow any reference count on the matched binding, while the (dmatch!) rule requires that the matched binding has a unique reference count:

$$\begin{array}{l} (\text{bmatch}_h) \quad H, y \mapsto^n C^k \bar{y} \mid \text{match } y \{ \overline{p \rightarrow e} \} \longrightarrow_h \quad H, y \mapsto^n C^k \bar{y} \mid e_i[\bar{x} := \bar{y}] \quad (p_i = C^k \bar{y}) \\ (\text{dmatch}_h) \quad H, x \mapsto^1 C^k \bar{y} \mid \text{match! } x \{ \overline{p \rightarrow e} \} \longrightarrow_h \quad H, \diamond_k \mid e_i[\bar{x} := \bar{y}] \quad (p_i = C^k \bar{x}) \end{array}$$

The extra transition rules are for general allocation and reference counting. The (alloc_h) rule allows allocating a constructor without a reuse credit, and similarly, the (lamb_h) rule allocates a closure. The application rule (app_h) applies a closure. We see that it starts by dupping its environment \bar{z} , and then dropping the closure itself. This way the APP rule can consider the free variables to part of the owned environment. This is important in practice as it allows a function to discard variables in the environment as soon as possible and be garbage-free. Here we can also see why we need to maintain \bar{z} explicitly: even though the free variables of a lambda expression are initially distinct, during evaluation some may be substituted by the same variable and we need to dup such variable multiple times when applying to maintain proper reference counts.

The other rules all deal with reference counting. The (dup_h) transition increments a reference count, while (drop_h) decrements a reference count $n > 1$. The (dlamb_h) rule drops a closure when the reference count is 1; this never creates a reuse credit though as the size of a closure cannot be accounted for statically. Note that also the environment is dropped, just like the (dconru_h) rule and the (dcon_s) rule in Figure 6. The (dconru_h) rule creates a reuse credit if the reference count is unique, while the (dropru_h) rule applies for a non-unique reference count with $n > 1$; this rule decrements the reference count but also allocates a fresh reuse credit (as required by rule DROPRU) – this is where the runtime falls back to copying if the cell was not unique.

Of course, with our match! translation, we no longer require the (dmatch_h) rule and can derive the rule from the translated expression when we assume the matched binding is unique:

$$\begin{array}{l} H, x \mapsto^1 C_i^k \bar{y} \mid \text{match! } x \{ C_i \bar{x}_i \rightarrow e_i \} \\ = \quad H, x \mapsto^1 C_i^k \bar{y} \mid \text{match } x \{ C_i \bar{x}_i \rightarrow \text{dup } \bar{x}_i; \text{dropru } x; e_i \} \\ \longrightarrow_h \quad H, x \mapsto^1 C_i^k \bar{y} \mid \text{dup } \bar{y}; \text{dropru } x; e_i[\bar{x}_i := \bar{y}] \\ \longrightarrow_h \quad H', x \mapsto^1 C_i^k \bar{y} \mid \text{dropru } x; e_i[\bar{x}_i := \bar{y}] \quad \{ H' \text{ is } H \text{ but with all } \bar{y} \text{ refcounts } +1 \text{ (A)} \} \\ \longrightarrow_h \quad H', \diamond_k \mid \text{drop } \bar{y}; e_i[\bar{x}_i := \bar{y}] \quad \{ (\text{dconru}_h), \text{rc is } 1 \} \\ \longrightarrow_h \quad H, \diamond_k \mid e_i[\bar{x}_i := \bar{y}] \quad \{ (\text{drop}_h), \text{(A)} \} \end{array}$$

However, the translation is more general, and can also proceed if the matched binding is not unique but shared – in that case the final steps use (dropru_h) and become:

$$\begin{array}{l} \dots \\ \longrightarrow_h \quad H', x \mapsto^{n+1} C_i^k \bar{y} \mid \text{dropru } x; e_i[\bar{x}_i := \bar{y}] \quad \{ H' \text{ is } H \text{ but with all } \bar{y} \text{ refcounts } +1 \text{ (A)} \} \\ \longrightarrow_h \quad H', \diamond_k, x \mapsto^n C_i^k \bar{y} \mid \text{drop } \bar{y}; e_i[\bar{x}_i := \bar{y}] \quad \{ (\text{dropru}_h) \} \\ \longrightarrow_h \quad H, \diamond_k, x \mapsto^n C_i^k \bar{y} \mid e_i[\bar{x}_i := \bar{y}] \quad \{ (\text{drop}_h), \text{(A)} \} \end{array}$$

where the binding for x stays alive but we still allocate a fresh reuse credit. This is exactly where we can generate the code shown in Section 1.4 where we essentially inline and specialize the definition of dropru and check upfront if the matched binding is unique or not.

5.3 Soundness of the Heap Semantics

First we generalize the properties of the store semantics to reference counted heaps:

Definition 3. (*Heap Soundness and Linearity*)

For a heap H we write $\text{dom}(H)$ to denote the set of variables x bound in H and write $\text{rng}(H)$ to denote the set of values $C \bar{x}$ bound in H . Two heaps H_1, H_2 are compatible if they map equal names $x \mapsto^n v \in H_1, x \mapsto^m w \in H_2$, to equal values $v = w$. A heap is *sound* if all free variables in

$\text{rng}(H)$ are bound: $\text{fv}(\text{rng}(H)) \subseteq \text{dom}(H)$. A heap is *linear* if it is sound, and any variable $x \mapsto^n v$ in $\text{dom}(H)$ occurs at most n times in the free variables of $\text{rng}(H)$. By $\text{roots}(H)$ we denote the multi-set of reuse credits of H and variables $x \mapsto^n v$ of $\text{dom}(H)$, which contains any variable $n - m$ times, if it occurs m times in the free variables of $\text{rng}(H)$.

The definition of linearity ensures that mutation is safe if the reference count is one. Exactly as in the store semantics, we write $[H]\bar{x}$ to denote a substitution that recursively replaces variables by their bound value in H . We assume that we are given heaps corresponding to the owned and borrowed values, but only require that the heap of the owned values is linear and do not assume that the heaps have a disjoint domain. Instead we use the *join* operator \otimes to define a joined heap of the borrowed and owned part, even if they have common elements with the same name and value. We join common elements by summing their reference counts. Since this eliminates one reference to their children, we decrease their reference count accordingly:

$$\begin{aligned} \emptyset \otimes H_2 &= H_2 \\ H_1, \diamond_k \otimes H_2 &= H_1 \otimes H_2, \diamond_k \\ H_1, x \mapsto^n v \otimes H_2 &= H_1 \otimes H_2, x \mapsto^n v \quad \text{iff } x \notin \text{dom}(H_2) \\ H_1, x \mapsto^n v \otimes H_2, x \mapsto^m v, z \mapsto^{k+1} w &= H_1 \otimes H_2, x \mapsto^{n+m} v, z \mapsto^k w \quad \text{iff } \bar{z} = \text{fv}(v) \end{aligned}$$

The rootset of $H_1 \otimes H_2$ is exactly the disjoint union of the roots of H_1 and H_2 . When applied to linear heaps, \otimes can be viewed as a partial commutative monoid [Jung et al. 2015] where every result is valid. However, we prefer a categorical view on linear heaps where a morphism $H_1 \rightarrow H_2$ exists if $\text{dom}(H_1) \subseteq \text{dom}(H_2)$ and $\text{roots}(H_1) \subseteq \text{roots}(H_2)$. This forms a monoidal category with the join operator as tensor product. We also define a heap subtraction operation $[H_1, H_2]$ if $H_1 \rightarrow H_2$ similar to an internal hom (and which corresponds to the magic wand of separation logic). We can then show that heap evaluation leaves the borrowed values unchanged:

Theorem 6. (The heap semantics is sound for well-formed Perceus programs)

If $\Delta \mid \Gamma \vdash e$ and given heaps H_1, H_2 with $\Delta \subseteq \text{dom}(H_1)$, H_1 sound, $\Gamma = \text{roots}(H_2)$ and H_2 linear, then $[H_1 \otimes H_2]e \mapsto^* \bar{v}$ implies $H_1 \otimes H_2 \mid e \mapsto_{\text{h}}^* H_1 \otimes H_3 \mid \bar{x}$ where $[H_3]\bar{x} = \bar{v}$, $\bar{x} = \text{roots}(H_3)$ and H_3 is linear.

This is again a strong theorem as it shows that the dynamic reference count is always correct and no variables will be discarded too early, while also having no garbage at the end of an evaluation ($\bar{x} = \text{roots}(H_3)$). Our proof (in Appendix E) is novel and may be well suited to possible mechanized formalization. As a corollary, any closed λ^{fip} expression can evaluate starting from an empty heap:

Corollary 3.

If $e \mapsto^* \bar{v}$ and $\emptyset \mid \emptyset \vdash e$, then $\emptyset \mid e \mapsto_{\text{h}}^* H \mid \bar{x}$ and $[H]\bar{x} = \bar{v}$.

While it is outside the scope of this paper, we could also modify the LET rule of our calculus with a (\star)-condition to characterize *garbage-free* and *frame-limited* derivations [Lorenzen and Leijen 2022]. However, borrowing makes it harder to achieve these properties and further study is needed. In particular, a garbage-free derivation can only exist if all borrowed arguments are still used later on, and similarly, a frame-limited derivation can only exist if all borrowed arguments are either used later on or have constant size.

6 BENCHMARKS

Figure 11 shows benchmark results of examples from this paper, relative to the *fip* variant. The results are the average over 5 runs on an AMD7950X on Ubuntu 22.04.2 with Koka v2.4.1-dev-fbip. Each benchmark uses 100 iterations over $N (=100000)$ element structures. We test each benchmark in following variants:

- *fip*: the algorithm implemented as FIP in Koka.

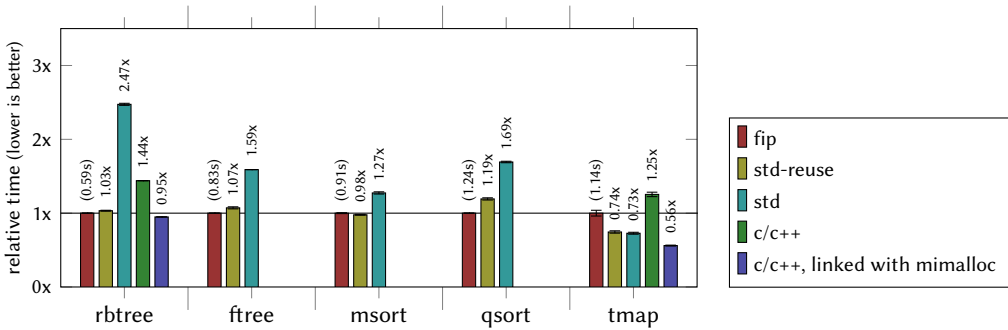


Fig. 11. Benchmarks on Ubuntu 22.04.2 (AMD 7950x), Koka v2.4.1-dev-fbip.

- *std*: the standard functional implementation in Koka without reuse optimization. For general GC'd languages without precise reference counts, the relative performance between *std* and *fip* can be more indicative of potential performance gains.
- *std-reuse*: just as *std* but with reuse optimization enabled. This is standard Koka which always applies dynamic reuse.
- *c/c++*: an standard in-place updating implementation in C or C++. Since our benchmarks are allocation heavy, we also include a variant when linked with the `mimalloc` [Leijen et al. 2019] memory allocator since that is usually faster than the standard C/C++ one.

The benchmarks consist of:

- *rbtree*: performs N balanced red-black tree insertions and folds the tree to compute the sum of the elements. The *fip* variant is the one in Section 4.1. while *std* uses Okasaki style insertion [Okasaki 1999]. The C++ versions use the standard in-place updating STL `std::map` which is implemented using red-black trees internally.
- *ftree*: builds a finger tree of size N and performs $3*N$ uncons/snoc operations. The *fip* variant is shown in Section 4.3 where the *std* variant uses a implementation described by Claessen [2020].
- *msort*, *qsort*: sorts an N element random list. The *fip* variant uses the implementations shown in Section 4.2 while *std* uses the standard recursive functional implementations (derived from the Haskell library implementations).
- *tmap*: maps an increment function over a *shared* (non-unique) N element tree returning a fresh tree which is then folded to compute the sum of the elements. The *fip* variant uses the implementation of Section 3 while *std* and *c/c++* use the standard (recursive) way to map over a tree.

It is hard to draw firm conclusions as the results are dependent on our particular implementation, but we make some general observations:

- The performance of *fip* versus *std* is generally much better showing that in-place updating is indeed generally faster than allocation.
- Even without a `fip` annotation, the *std-reuse* variant shows that the reuse optimization in Koka can be very effective – but of course, unlike `fip`, reuse here is not guaranteed.
- In an absolute sense, the performance seems very good where in the *rbtree* benchmark the *fip* variant rivals the performance of the in-place updating `std::map` implementation in C++.
- The *tmap* benchmark is interesting as *fip* is generally slower here. The *fip* variant uses a zipper to visit the tree such that uses constant stack space (unlike the others which use stack space linear in the depth of the tree). Reversing the pointers Schorr-Waite style can be slower though than recursing with the stack. Also, the tree that is mapped is shared and thus even the `fip` function

FP²: Fully in-Place Functional Programming

cannot reuse the original tree. Nevertheless, the *fip* variant will still reuse the zipper it uses to traverse the tree. This also shows why *std* and *std-reuse* are performing similarly since there is no reuse possible for the standard algorithms. That *std-reuse* is only about 1% slower shows that the dynamic reuse check has negligible impact on performance.

7 RELATED WORK

The FIP calculus is most closely related to Hofmann’s type system for in-place update [Hofmann 2000b 2000a]. Just like Hofmann, we add reuse credits to a linear environment, model a destructive match, and collect top-level functions in the signature. However, Hofmann’s unboxed tuples can escape into allocations, which makes it necessary to monomorphise the program (and track types to be able to do so). In contrast, our calculus does not need monomorphisation or know about types at all. Hofmann also uses a uniform size for all constructors of a datatype (including atoms such as `Nil`), but unboxes the first layer of each datatype. Many FIP programs can also be checked by that scheme, but it seems to increase memory usage substantially: in our calculus, a constructor with n fields filled with atoms takes n space, while it would take $n * n$ space in Hofmann’s calculus.

While we only model unique and borrowed values in our FIP calculus, *shared* values are another interesting variant. Unlike borrowed values, shared values can be stored in datatypes. But unlike unique values, they can be used multiple times (and it is not possible to use a destructive `match!` on them). Shared values correspond to the usage aspect 2 introduced by Aspinall and Hofmann [2002] and Aspinall et al. [2008]. We believe that it may be worthwhile to extend the FIP calculus with shared values to allow it to check a wider range of programs. However, shared values can only be supported in a garbage-collected setting, while our FIP programs can also easily be compiled to C.

Even without in-place reuse, FIP programs still use constant space, which allows us to reason about their space usage. Space credits [Hofmann 2003; Hofmann and Jost 2003] generalize reuse credits with the axiom $\diamond_{n_1}, \diamond_{n_2} = \diamond_{n_1 + n_2}$. This axiom does not hold for reuse credits (which can not be combined unless they are in adjacent slots in the heap), but it does hold if we view \diamond_n just as the promise that n words of space is available. Based on space credits, an automated analysis [Hofmann et al. 2011; Hofmann and Jost 2006] or manual proofs in separation logic [Madiot and Pottier 2022; Moine et al. 2023] can be used to reason about heap space. However, these systems usually do not model atoms or unboxing, which we identified as crucial for real-world FIP programs.

Reuse analysis can be implemented either statically using uniqueness types [Barendsen and Smetsers 1995] or flow analysis, or dynamically using reference counts. Compile-time Garbage Collection [Bruynooghe 1986] is the most developed flow based analysis, which tracks the flows of unique values through the program to identify reuse opportunities statically. Reuse with reference counts has long been applied to arrays, where the update function can be designed to mutate the array in-place if the reference count is one [Hudak and Bloss 1985; Scholz 1994]. Similarly, Stoye et al. [1984] used reuse with one-bit reference counts for combinator reduction, where reuse is encoded in the hand-written combinators. However, in this work, we rely on a reuse analysis that can statically discover reuse opportunities between otherwise unconnected memory cells, which was pioneered by OPAL [Didrich et al. 1994; Schulte 1994; Schulte and Grieskamp 1992]. Their analysis was refined by Ullrich and de Moura [2019], who showed that such an analysis can be implemented efficiently without duplicating code. Reinking, Xie et al. [2021] present the linear resource calculus as a formalization of precise reference counting and give a *garbage free* algorithm. Lorenzen and Leijen [2022] refine this calculus further with a declarative star condition that can guarantee either garbage-free or *frame-limited* space usage which ensures extra space usage due to reuse is bounded.

Borrowing is a long-standing technique to reduce the overhead of reference counting [Baker 1994b; Lemaitre et al. 1986]. If a lifetime analysis can prove that the lifetime of one reference dominates

that of another, we can avoid counting the second reference – in our calculus, this is expressed by borrowing Γ_2 in `LET` (first introduced by [Reinking, Xie et al. 2021]). Ullrich and de Moura [2019] introduced borrowed parameters on top-level functions which is especially important for recursive functional code. They also showed an inference for borrowing annotations, but, as pointed out by Lorenzen and Leijen [2022], this can increase memory usage by an unbounded amount.

Wadler [Wadler 1984] describes the *listless* transformation that can transform compositions of many common list functions, like `map`, `append`, or `partition`, into a state machine that does not allocate any `Cons` cells. Functions that cannot be transformed include our `reverse` example, sorting a list, or handling trees. Since the transform can discard `Cons` cells, we believe that any listless program is also a well-formed FBIP program.

8 CONCLUSION AND FUTURE WORK

We have shown the necessary features for, and spirit of, fully in-place functional programming. We believe the examples given in this paper have only scratched the surface of what is possible and that there are more FIP algorithms waiting to be discovered. Another interesting research direction is to extend the fully in-place calculus to cover more possible programs that are currently not quite FIP; for example by adding shared values as described by Aspinall et al. [2008].

REFERENCES

- David Aspinall, and Martin Hofmann. 2002. Another Type System for in-Place Update. In *ESOP*, 2:36–52. Springer.
- David Aspinall, Martin Hofmann, and Michal Konečný. 2008. A Type System with Usage Aspects. *Journal of Functional Programming* 18 (2). Cambridge University Press: 141–178.
- Roland C Backhouse. 1988. *An Exploration of the Bird-Meertens Formalism*. University of Groningen, Department of Mathematics and Computing Science.
- Henry G Baker. 1994a. A “linear Logic” Quicksort. *ACM Sigplan Notices* 29 (2). ACM New York, NY, USA: 13–18.
- Henry G Baker. 1994b. Minimizing Reference Count Updating with Deferred and Anchored Pointers for Functional Data Structures. *ACM Sigplan Notices* 29 (9). ACM New York, NY, USA: 38–43.
- Erik Barendsen, and Sjaak Smetsers. 1995. Uniqueness Type Inference. In *Proceedings of the 7th International Symposium on Programming Languages: Implementations, Logics and Programs*, 189–206.
- Frédéric Bour, Basile Clément, and Gabriel Scherer. Apr. 2021. Tail Modulo Cons. *Journées Francophones Des Langages Applicatifs (JFLA)*, April. Saint Médard d’Excideuil, France. <https://hal.inria.fr/hal-03146495/document>. hal-03146495.
- Tom H Brus, Marko CJD van Eekelen, MO Van Leer, and Marinus J Plasmeijer. 1987. Clean—a Language for Functional Graph Rewriting. In *Functional Programming Languages and Computer Architecture: Portland, Oregon, USA, September 14–16, 1987 Proceedings*, 364–384. Springer.
- Maurice Bruynooghe. 1986. *Compile Time Garbage Collection*. Katholieke Universiteit Leuven. Departement Computerwetenschappen.
- Koen Claessen. 2020. Finger Trees Explained Anew, and Slightly Simplified (functional Pearl). In *Proceedings of the 13th ACM SIGPLAN International Symposium on Haskell*, 31–38.
- Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. 2022. *Introduction to Algorithms*. MIT press.
- Olivier Danvy. 2008. From Reduction-Based to Reduction-Free Normalization. In *Proceedings of the 6th International Conference on Advanced Functional Programming*, 66–164. AFP’08. Springer-Verlag, Berlin, Heidelberg.
- Olivier Danvy. 2022. Getting There and Back Again. *Fundamenta Informaticae* 185. Episciences. org.
- Edsko De Vries, Rinus Plasmeijer, and David M Abrahamson. 2008. Uniqueness Typing Simplified. In *Implementation and Application of Functional Languages: 19th International Workshop, IFL 2007, Freiburg, Germany, September 27-29, 2007. Revised Selected Papers* 19, 201–218. Springer.
- Klaus Didrich, Andreas Fett, Carola Gerke, Wolfgang Grieskamp, and Peter Pepper. 1994. OPAL: Design and Implementation of an Algebraic Programming Language. In *Programming Languages and System Architectures*, 228–244. Springer.
- Cormac Flanagan, Amr Sabry, Bruce F Duba, and Matthias Felleisen. 1993. The Essence of Compiling with Continuations. In *Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation*, 237–247.
- Jeremy Gibbons. 1994. An Introduction to the Bird- Meertens Formalism.
- J.R. Hindley. Dec. 1969. The Principal Type Scheme of an Object in Combinatory Logic. *Transactions of the American Mathematical Society* 146 (December): 29–60.
- Ralf Hinze, and Ross Paterson. 2006. Finger Trees: A Simple General-Purpose Data Structure. *Journal of Functional Programming* 16 (2). Cambridge University Press: 197–217.
- Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. 2011. Multivariate Amortized Resource Analysis. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 357–370.
- Martin Hofmann. 2000a. In-Place Update with Linear Types or How to Compile Functional Programms into Malloc-Free C. *Preprint, Www.Dcs.Ed.Ac.Uk/~Mxh/malloc.Ps.Gz*. Citeseer.
- Martin Hofmann. 2000b. A Type System for Bounded Space and Functional in-Place Update. In *European Symposium on Programming*, 165–179. Springer.
- Martin Hofmann. 2003. Linear Types and Non-Size-Increasing Polynomial Time Computation. *Information and Computation* 183 (1). Elsevier: 57–85.
- Martin Hofmann, and Steffen Jost. 2003. Static Prediction of Heap Space Usage for First-Order Functional Programs. *ACM SIGPLAN Notices* 38 (1). ACM New York, NY, USA: 185–197.
- Martin Hofmann, and Steffen Jost. 2006. Type-Based Amortised Heap-Space Analysis. In *European Symposium on Programming*, 22–37. Springer.
- Paul Hudak. 1986. A Semantic Model of Reference Counting and Its Abstraction (detailed Summary). In *Proceedings of the 1986 ACM Conference on LISP and Functional Programming*, 351–363.
- Paul Hudak, and Adrienne Bloss. 1985. The Aggregate Update Problem in Functional Programming Systems. In *Proceedings of the 12th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, 300–314.
- Gérard Huet. 1997. The Zipper. *Journal of Functional Programming* 7 (5). Cambridge University Press: 549–554.
- R John Muir Hughes. 1986. A Novel Representation of Lists and Its Application to the Function “reverse.” *Information Processing Letters* 22 (3). Elsevier: 141–144.
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. *ACM SIGPLAN Notices* 50 (1). ACM New York,

- NY, USA: 637–650.
- Daan Leijen, and Anton Lorenzen. 2023. Tail Recursion Modulo Context: An Equational Approach. *Proceedings of the ACM on Programming Languages* 7 (POPL). ACM New York, NY, USA: 1152–1181.
- Daan Leijen, Benjamin Zorn, and Leonardo de Moura. 2019. Mimalloc: Free List Sharding in Action. In *Asian Symposium on Programming Languages and Systems*, 244–265. Springer.
- Michel Lemaitre, Michel Castan, M-H Durand, Guy Durrieu, and Bernard Lecussan. 1986. Mechanisms for Efficient Multiprocessor Combinator Reduction. In *Proceedings of the 1986 ACM Conference on LISP and Functional Programming*, 113–121.
- Alexey Loginov, Thomas Reps, and Mooly Sagiv. 2006. Automated Verification of the Deutsch-Schorr-Waite Tree-Traversal Algorithm. In *Static Analysis: 13th International Symposium, SAS 2006, Seoul, Korea, August 29-31, 2006. Proceedings 13*, 261–279. Springer.
- Anton Lorenzen, and Daan Leijen. 2022. Reference Counting with Frame Limited Reuse. *Proceedings of the ACM on Programming Languages* 6 (ICFP). ACM New York, NY, USA: 357–380.
- Jean-Marie Madiot, and François Pottier. 2022. A Separation Logic for Heap Space under Garbage Collection. *Proceedings of the ACM on Programming Languages* 6 (POPL). ACM New York, NY, USA: 1–28.
- Conor McBride. 2001. The Derivative of a Regular Type Is Its Type of One-Hole Contexts. *Unpublished Manuscript*, 74–88.
- Conor McBride. 2008. Clowns to the Left of Me, Jokers to the Right (pearl) Dissecting Data Structures. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 287–295.
- Robin Milner. 1978. A Theory of Type Polymorphism in Programming. *Journal of Computer and System Sciences* 17: 248–375.
- Alexandre Moine, Arthur Charguéraud, and François Pottier. 2023. A High-Level Separation Logic for Heap Space under Garbage Collection. *Proceedings of the ACM on Programming Languages* 7 (POPL). ACM New York, NY, USA: 718–747.
- Chris Okasaki. 1999. Red-Black Trees in a Functional Setting. *Journal of Functional Programming* 9 (4). Cambridge University Press: 471–477.
- Simon L. Peyton Jones, and John Launchbury. 1991. Unboxed Values as First Class Citizens in a Non-Strict Functional Language. In *Functional Programming Languages and Computer Architecture*, edited by John Hughes, 636–666. Springer Berlin Heidelberg.
- Reinking, Xie, de Moura, and Leijen. 2021. Perceus: Garbage Free Reference Counting with Reuse. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 96–111. PLDI 2021. Virtual, Canada. doi:10.1145/3453483.3454032.
- John C Reynolds. 1972. Definitional Interpreters for Higher-Order Programming Languages. In *Proceedings of the ACM Annual Conference-Volume 2*, 717–740.
- John C Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science*, 55–74. IEEE.
- Sven-Bodo Scholz. 1994. Single Assignment C-Functional Programming Using Imperative Style. In *Proceedings of IFL*, volume 94.
- Herbert Schorr, and William M Waite. 1967. An Efficient Machine-Independent Procedure for Garbage Collection in Various List Structures. *Communications of the ACM* 10 (8). ACM New York, NY, USA: 501–506.
- Wolfram Schulte. 1994. Deriving Residual Reference Count Garbage Collectors. In *International Symposium on Programming Language Implementation and Logic Programming*, 102–116. Springer.
- Wolfram Schulte, and Wolfgang Grieskamp. 1992. Generating Efficient Portable Code for a Strict Applicative Language. In *Declarative Programming, Sasbachwalden 1991*, 239–252. Springer.
- Daniel Dominic Sleator, and Robert Endre Tarjan. 1985. Self-Adjusting Binary Search Trees. *Journal of the ACM* 32: 652–686.
- Jonathan Sobel, and Daniel P Friedman. 1998. Recycling Continuations. In *Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming*, 251–260.
- William R Stoye, Thomas JW Clarke, and Arthur C Norman. 1984. Some Practical Methods for Rapid Combinator Reduction. In *Proceedings of the 1984 ACM Symposium on LISP and Functional Programming*, 159–166.
- Sebastian Ullrich, and Leonardo de Moura. 2019. Counting Immutable Beans: Reference Counting Optimized for Purely Functional Programming. In *Proceedings of the 31st Symposium on Implementation and Application of Functional Languages*, 1–12.
- Twan van Laarhoven. 2007. Deriving Functor. <https://mail.haskell.org/pipermail/haskell-prime/2007-March/002137.html>.
- Philip Wadler. 1984. Listlessness Is Better than Laziness: Lazy Evaluation and Garbage Collection at Compile-Time. In *Proceedings of the 1984 ACM Symposium on LISP and Functional Programming*, 45–52. LFP’84. Austin, Texas, USA. doi:10.1145/800055.802020.
- David Walker, and Greg Morrisett. 2000. Alias Types for Recursive Data Structures. *Types in Compilation* 2071. Springer: 177–206.
- Hongseok Yang. 2007. Relational Separation Logic. *Theoretical Computer Science* 375 (1-3). Elsevier: 308–334.

A EXAMPLE CODE

This section contains the full examples discussed in Section 4. These examples do not use `match!`, `drop` or `free` as they are automatically inferred by our implementation. We also write `val (C x1 ... xn) = x in e` for `match x { C x1 ... xn -> e }` if the match has just a single case.

A.1 Red-Black Trees

```
import std/num/int32

type color
  Red
  Black

type tree
  Node(color : color, lchild : tree, key : int32, value : bool, rchild : tree)
  Leaf

fip fun is-red(^t : tree) : bool
  match t
    Node(Red) -> True
    _         -> False

type accum
  Done
  NodeL(color : color, lchild : accum, key : int32, value : bool, rchild : tree)
  NodeR(color : color, lchild : tree, key : int32, value : bool, rchild : accum)

fip(1) fun ins(t : tree, key : int32, v : bool, z : accum) : exn tree
  match t
    Node(c, l, kx, vx, r)
      -> if key < kx then ins(l, key, v, NodeL(c, z, kx, vx, r))
        elif key > kx then ins(r, key, v, NodeR(c, l, kx, vx, z))
        else balance(z, Node(c, l, key, v, r))
    Leaf -> balance(z, Node(Red, Leaf, key, v, Leaf))

fip fun set-black(t : tree) : tree
  match t
    Node(_, l, k, v, r) -> Node(Black, l, k, v, r)
    t -> t

fip fun rebuild(z : accum, t : tree) // Turn the zipper into a tree without rotating
  match z
    NodeR(c, l, k, v, z1) -> rebuild(z1, Node(c, l, k, v, t))
    NodeL(c, z1, k, v, r) -> rebuild(z1, Node(c, t, k, v, r))
    Done -> t

fip(1) fun insert(t : tree, k : int32, v : bool) : <exn> tree
  ins(t, k, v, Done)

fip fun balance( z : accum, t : tree ) : exn tree
  match z
    NodeR(Red, l1, k1, v1, z1) -> match z1
      NodeR(_,l2,k2,v2,z2) -> // black
        if is-red(l2) then balance(z2, Node(Red, l2.set-black, k2, v2, Node(Black, l1, k1, v1, t) ))
        else rebuild(z2, Node(Black, Node(Red,l2,k2,v2,l1), k1, v1, t))
      NodeL(_,z2,k2,v2,r2) -> // black
        if is-red(r2) then balance(z2, Node(Red, Node(Black,l1,k1,v1,t), k2, v2, r2.set-black))
        else match t
          Node(_, l, k, v, r) ->
            rebuild(z2, Node(Black, Node(Red,l1,k1,v1,l), k, v, Node(Red,r,k2,v2,r2)))
          Done -> Node(Black, l1, k1, v1, t)
    NodeL(Red, z1, k1, v1, r1) -> match z1
      NodeL(_,z2,k2,v2,r2) -> // black
        if is-red(r2) then balance(z2, Node(Red, Node(Black, t, k1, v1, r1), k2, v2, r2.set-black ))
        else rebuild(z2, Node(Black, t, k1, v1, Node(Red,r1,k2,v2,r2)))
```



```

NodeR(_,l2,k2,v2,z2) -> // black
  if is-red(l2) then balance(z2, Node(Red, l2.set-black, k2, v2, Node(Black,t,k1,v1,r1) ))
  else match t
    Node(_, l, k, v, r) ->
      rebuild(z2, Node(Black, Node(Red,l2,k2,v2,l), k, v, Node(Red,r,k1,v1,r1)))
  Done -> Node(Black, t, k1, v1, r1)
z -> rebuild(z, t)

```

A.2 Stable Merge Sort

Derived from the merge sort of Haskell's `Data.List.sort` function.

```

import std/num/int32
import std/os/env

alias elem = int32

ref type pad
  Pad

type unit2
  Unit2(a : pad, b : pad)

type pair<a>
  Pair(a : a, b : a)

type sublist<a>
  SCons(a : a, cs : sublist<a>)
  STuple(a : a, b : a)

type partition<a>
  Sublist(c : sublist<a>, z : partition<a>)
  Singleton(c : a, z : partition<a>)
  End

fip fun reverse-go(c : sublist<a>, acc : sublist<a>, u : unit2) : sublist<a>
  match c
    SCons(a, cs) -> reverse-go(cs, SCons(a, acc), u)
    STuple(a, b) -> SCons(b, SCons(a, acc))

fip fun reverse-sublist(c : sublist<a>) : sublist<a>
  match c
    SCons(a, SCons(b, c)) -> reverse-go(c, STuple(b, a), Unit2(Pad,Pad))
    SCons(a, STuple(b, c)) -> SCons(c, STuple(b, a))
    STuple(a, b) -> STuple(b, a)

fip fun sequences(xs : list<elem>) : div partition<elem>
  match(xs)
    Cons(a, Cons(b, xs1)) -> if(a > b)
      then
        val (sublist, bs) = descending(b, STuple(b, a), xs1)
        Sublist(sublist, sequences(bs))
      else
        val (sublist, bs) = ascending(b, STuple(b, a), xs1)
        Sublist(sublist, sequences(bs))
    Cons(a, Nil) -> Singleton(a, End)
    Nil -> End

fip fun descending(a : elem, sublist : sublist<elem>, bs : list<elem>) : (sublist<elem>, list<elem>)
  match(bs)
    Cons(b, bs1) | a > b -> descending(b, SCons(b, sublist), bs1)
    bs -> (sublist, bs)

```

FP²: Fully in-Place Functional Programming

```

fix fun ascending(a : elem, sublist : sublist<elem>, bs : list<elem>) : (sublist<elem>, list<elem>)
  match (bs)
    Cons(b, bs1) | (a <= b) -> ascending(b, SCons(b, sublist), bs1)
    bs -> (reverse-sublist(sublist), bs)

fix fun to-list(c : sublist<a>, u : unit2) : list<a>
  match c
    SCons(a, cs) -> Cons(a, to-list(cs, u))
    STuple(a, b) -> Cons(a, Cons(b, Nil))

fix fun merge-all(xs : partition<elem>) : <div> list<elem>
  match (xs)
    Sublist(x, End) -> to-list(x, Unit2(Pad, Pad))
    Singleton(x, End) -> Cons(x, Nil)
    xs -> merge-all(merge-pairs(xs))

fix fun merge-pairs(xs : partition<elem>) : <div> partition<elem>
  match (xs)
    Sublist(a, Sublist(b, xs1)) -> Sublist(merge(a, b, Unit2(Pad, Pad)), merge-pairs(xs1))
    Sublist(a, Singleton(b, xs1)) -> Sublist(merge-last-left(a, b, Unit2(Pad, Pad)), merge-pairs(xs1))
    Singleton(a, Sublist(b, xs1)) -> Sublist(merge-last-right(a, b, Unit2(Pad, Pad)), merge-pairs(xs1))
    Singleton(a, Singleton(b, xs1)) ->
      Sublist(if a <= b then STuple(a, b) else STuple(b, a), merge-pairs(xs1))
    xs -> xs

fix fun merge(c1 : sublist<elem>, c2 : sublist<elem>, u : unit2) : <div> sublist<elem>
  match c1
    SCons(a, cs1) -> match c2
      SCons(b, cs2) ->
        if a <= b then SCons(a, merge(cs1, SCons(b, cs2), u))
        else SCons(b, merge(SCons(a, cs1), cs2, u))
      STuple(b, c) ->
        if a <= b then SCons(a, merge(cs1, STuple(b, c), u))
        else SCons(b, merge-last-left(SCons(a, cs1), c, u))
    STuple(a, b) -> match c2
      SCons(c, cs2) ->
        if a <= c then SCons(a, merge-last-right(b, SCons(c, cs2), u))
        else SCons(c, merge(STuple(a, b), cs2, u))
      STuple(c, d) ->
        if a <= c then SCons(a, merge-right(b, Pair(c, d), u))
        else SCons(c, merge-left(Pair(a, b), d, u))

fix fun merge-last-right(a : elem, c2 : sublist<elem>, u : unit2) : sublist<elem>
  match c2
    SCons(b, cs2) -> if a <= b
      then SCons(a, SCons(b, cs2))
      else SCons(b, merge-last-right(a, cs2, u))
    STuple(b, c) -> merge-right(a, Pair(b, c), u)

fix fun merge-last-left(c2 : sublist<elem>, d : elem, u : unit2) : sublist<elem>
  match c2
    SCons(a, cs2) -> if a <= d
      then SCons(a, merge-last-left(cs2, d, u))
      else SCons(d, SCons(a, cs2))
    STuple(a, b) -> merge-left(Pair(a, b), d, u)

fix fun merge-right(a : elem, p : pair<elem>, u : unit2) : sublist<elem>
  match p
    Pair(b, c) -> if a <= b
      then SCons(a, STuple(b, c))
      else SCons(b, if a <= c then STuple(a, c) else STuple(c, a))

fix fun merge-left(p : pair<elem>, d : elem, u : unit2) : sublist<elem>
  match p
    Pair(a, b) -> if a <= d
      then SCons(a, if b <= d then STuple(b, d) else STuple(d, b))
      else SCons(d, STuple(a, b))

```

A.3 Quick Sort

```

import std/num/int32
import std/os/env

alias elem = int32

ref type pad
  Pad

ref type unit2
  Unit2(a : pad, b : pad)

type maybe2<a>
  Nothing2
  Just2(a : a, b : pad)

type sublist<a>
  SCons(a : a, cs : sublist<a>)
  STuple(a : a, b : a)

type partition<a>
  Sublist(c : sublist<a>, bdl : partition<a>)
  Singleton(c : a, bdl : partition<a>)
  End

fip fun from-list(x, y, u : unit2, xs : list<elem>) : sublist<elem>
  match xs
  Cons(z, zs) -> SCons(x, from-list(y, z, u, zs))
  Nil -> STuple(x, y)

fip fun to-bundle(xs : list<elem>, bdl' : partition<elem>) : partition<elem>
  match xs
  Cons(x, Cons(y, xs)) -> Sublist(from-list(x, y, Unit2(Pad, Pad), xs), bdl')
  Cons(x, Nil) -> Singleton(x, bdl')
  Nil -> bdl'

fip fun quicksort(xs : list<elem>) : div list<elem>
  quicksort-go(to-bundle(xs, End))

fip fun quicksort-go(bdl : partition<elem>) : div list<elem>
  match bdl
  Sublist(xs, bdl') -> match xs
  SCons(p, xs) ->
    val (lo, hi) = partition(p, xs, Unit2(Pad, Pad))
    quicksort-go(to-bundle(lo, Singleton(p, to-bundle(hi, bdl'))))
  STuple(x, y) | x <= y -> Cons(x, Cons(y, quicksort-go(bdl')))
  | _ -> Cons(y, Cons(x, quicksort-go(bdl')))
  Singleton(p, b) -> Cons(p, quicksort-go(b))
  End -> Nil

fip fun partition(^p : elem, xs : sublist<elem>, u : unit2) : (list<elem>, list<elem>)
  match(xs)
  SCons(x, xx) -> if p <= x
  then val (lo, hi) = partition(p, xx, u)
  (lo, Cons(x, hi))
  else val (lo, hi) = partition(p, xx, u)
  (Cons(x, lo), hi)
  STuple(x, y) -> if p <= x
  then if p <= y
  then (Nil, Cons(x, Cons(y, Nil)))
  else (Cons(y, Nil), Cons(x, Nil))
  else if p <= y
  then (Cons(x, Nil), Cons(y, Nil))
  else (Cons(x, Cons(y, Nil)), Nil)

```

FP²: Fully in-Place Functional Programming

A.4 Finger Trees

```
// Adapted from "Finger Trees Explained Anew, and Slightly Simplified (Functional Pearl)", Claessen
import std/num/int32

ref type pad
  Pad

type reuse3
  Reuse3(a : pad, b : pad, c : pad)

type afew<a>
  One(a : a, b : pad, c : pad)
  Two(a : a, b : a, c : pad)
  Three(a : a, b : a, c : a)

type tuple<a>
  Pair(a : a, b : a, c : pad)
  Triple(a : a, b : a, c : a)

type seq<a>
  Empty
  Unit(a : a, b : pad, c : pad)
  More0(l : a, s : seq<tuple<a>>, r : afew<a>)
  More(l : tuple<a>, s : seq<tuple<a>>, r : afew<a>)

type buffer
  BNil
  BCons(next : buffer, b : pad, c : pad)

value type bseq<a>
  BSeq(s : seq<a>, q : buffer)

// Isomorphic to (,,) but unboxed
value type tuple4<a,b,c,d>
  Tuple4(fst:a,snd:b,thd:c,field4:d)

fun bhead(^bs : bseq<a>) : exn a
  match bs
    BSeq(s, _) -> head(s)

fun head(^s : seq<a>) : exn a
  match s
    Unit(x) -> x
    More0(x, _, _) -> x
    More(Pair(x, _, _), _, _) -> x
    More(Triple(x, _, _), _, _) -> x

fip fun bcons(x : a, u3 : reuse3, bs : bseq<a>) : exn bseq<a>
  val BSeq(s, b) = bs
  val (s', b') = cons(x, u3, s, b)
  BSeq(s', b')

fip fun cons(x : a, u3 : reuse3, s : seq<a>, b : buffer) : exn (seq<a>, buffer)
  match s
    Empty -> (Unit(x, Pad, Pad), b)
    Unit(y, _, _) -> (More0(x, Empty, One(y, Pad, Pad)), b)
    More0(y, q, u) -> (More(Pair(x, y, Pad), q, u), b)
    More(Pair(y, z, _), q, u) ->
      (More(Triple(x, y, z), q, u), BCons(b, Pad, Pad))
    More(Triple(y, z, w), q, u) ->
      val BCons(b', _, _) = b
      val (q', b'') = cons(Pair(z, w, Pad), u3, q, b')
      (More(Pair(x, y, Pad), q', u), b'')

fip fun buncons(bs : bseq<a>) : exn (a, reuse3, bseq<a>)
  val BSeq(s, b) = bs
  val Tuple4(x, u3, s', b') = uncons(s, b)
  (x, u3, BSeq(s', b'))
```

```

fix fun uncons(s : seq<a>, b : buffer) : exn tuple4<a, reuse3, seq<a>, buffer>
  match s
  Unit(x, _, _) ->
    Tuple4(x, Reuse3(Pad,Pad,Pad), Empty, b)
  More(Triple(x, y, z), q, u) ->
    val BCons(b', _, _) = b
    Tuple4(x, Reuse3(Pad,Pad,Pad), More(Pair(y, z, Pad), q, u), b')
  More(Pair(x, y, _) , q, u) ->
    Tuple4(x, Reuse3(Pad,Pad,Pad), More0(y, q, u), b)
  More0(x, q, u) ->
    val (q', b') = more0(q, u, b)
    Tuple4(x, Reuse3(Pad,Pad,Pad), q', b')

fix fun more0(q : seq<tuple<a>>, u : afew<a>, b : buffer) : exn (seq<a>, buffer)
  match q
  Empty ->
    match u
    One(x, y, z) -> (Unit(x, y, z), b)
    Two(y, z, _) ->
      val BCons(b', _, _) = b
      (More0(y, Empty, One(z, Pad, Pad)), b')
    Three(y, z, w) ->
      val BCons(b', _, _) = b
      (More0(y, Empty, Two(z, w, Pad)), b')
  Unit(p, _, _) ->
    match p
    Pair(x, y, _) -> (More(Pair(x, y, Pad), Empty, u), b)
    Triple(x, y, z) ->
      val BCons(b', _, _) = b
      (More0(x, Unit(Pair(y,z,Pad),Pad,Pad), u), b')
  More0(p, q1, u1) ->
    match p
    Pair(x, y) ->
      val (q1', b') = more0(q1, u1, b)
      (More(Pair(x, y, Pad), q1', u), b')
    Triple(x, y, z) ->
      val BCons(b', _, _) = b
      (More0(x, More0(Pair(y,z,Pad), q1, u1), u), b')
  More(Pair(p, y1), q1, u1) ->
    match p
    Pair(x, y) -> (More(Pair(x, y, Pad), More0(y1, q1, u1), u), b)
    Triple(x, y, z) ->
      val BCons(b', _, _) = b
      (More0(x, More(Pair(Pair(y,z,Pad), y1, Pad), q1, u1), u), b')
  More(Triple(p, y1, z1), q1, u1) ->
    val BCons(b', _, _) = b
    match p
    Pair(x, y) ->
      (More(Pair(x, y, Pad), More(Pair(y1, z1, Pad), q1, u1), u), b')
    Triple(x, y, z) ->
      (More0(x, More(Triple(Pair(y,z,Pad), y1, z1), q1, u1), u), b')

fix fun bsnoc(bs : bseq<a>, u3 : reuse3, x : a) : exn bseq<a>
  val BSeq(s, b) = bs
  val (s', b') = snoc(s, b, u3, x)
  BSeq(s', b')

```

FP²: Fully in-Place Functional Programming

```

fip fun snoc(s : seq<a>, b : buffer, u3 : reuse3, x : a) : exn (seq<a>, buffer)
  match s
  Empty                               -> (Unit(x, Pad, Pad), b)
  Unit(y, _, _)                       -> (More0(y, Empty, One(x, Pad, Pad)), b)
  More0(u, q, One(y, _, _))           -> (More0(u, q, Two(y, x, Pad)), BCons(b, Pad, Pad))
  More (u, q, One(y, _, _))           -> (More (u, q, Two(y, x, Pad)), BCons(b, Pad, Pad))
  More0(u, q, Two(y, z, _))           -> (More0(u, q, Three(y, z, x)), BCons(b, Pad, Pad))
  More (u, q, Two(y, z, _))           -> (More (u, q, Three(y, z, x)), BCons(b, Pad, Pad))
  More0(u, q, Three(y, z, w)) ->
    val BCons(b', _, _) = b
    val (q', b'') = snoc(q, b', u3, Pair(y, z, Pad))
    (More0(u, q', Two(w, x, Pad)), b'')
  More(u, q, Three(y, z, w)) ->
    val BCons(b', _, _) = b
    val (q', b'') = snoc(q, b', u3, Pair(y, z, Pad))
    (More(u, q', Two(w, x, Pad)), b'')

```

B SOUNDNESS OF STORE SEMANTICS

This section contains the soundness result for the store semantics. It follows a typical progress-and-preservation style proof, where we maintain an invariant across all steps of the store semantics.

First, we define a simple invariant and show that this guarantees progress of evaluation if we omit the `eval` rule. We also show that our invariant is maintained by such evaluation steps. But this is not quite enough since the `eval` rule has a complicated interaction with the `bmatch` construct. We explain how our simple invariant goes wrong and define a more complicated one, which is also preserved by the `eval` rule. Together, these results directly imply the soundness claim.

We define the free and bound variables in the usual way. In this section, we omit constructors and top-level, statically-called functions from the free variables.

$\text{fv}(f)$	$:= \emptyset$
$\text{fv}(x)$	$:= \{x\}$
$\text{fv}(C^k v_1 \dots v_k)$	$:= \text{fv}(v_1), \dots, \text{fv}(v_k)$
$\text{fv}((v_1, \dots, v_n))$	$:= \text{fv}(v_1), \dots, \text{fv}(v_n)$
$\text{fv}(e_1 e_2)$	$:= \text{fv}(e_1), \text{fv}(e_2)$
$\text{fv}(f(e_1; e_2))$	$:= \text{fv}(e_1), \text{fv}(e_2)$
$\text{fv}(\text{let } \bar{x} = e_1 \text{ in } e_2)$	$:= \text{fv}(e_1), \text{fv}(e_2) - \bar{x}$
$\text{fv}(\text{match } e \{ \bar{p} \mapsto \bar{e} \})$	$:= \text{fv}(e), \text{fv}(e_1) - \text{bv}(p_1), \dots, \text{fv}(e_n) - \text{bv}(p_n)$
$\text{fv}(\text{match}! e \{ \bar{p} \mapsto \bar{e} \})$	$:= \text{fv}(e), \text{fv}(e_1) - \text{bv}(p_1), \dots, \text{fv}(e_n) - \text{bv}(p_n)$
$\text{bv}(C^k x_1 \dots x_k)$	$:= \{x_1, \dots, x_k\}$
$\text{fv}(\text{drop } \bar{x}; e)$	$:= \text{fv}(e), \bar{x}$
$\text{fv}(\text{free } k; e)$	$:= \text{fv}(e)$

B.1 Properties of Stores

In the following, we will often need to focus on a specific part of the store that corresponds to all values reachable from a root set Γ . For a linear store with $\Gamma \subseteq \text{roots}(S)$, we write $S[\Gamma]$ for the smallest linear subset of S containing Γ . Then we can split any store with roots Γ_1, Γ_2 into $S[\Gamma_1]$, $S[\Gamma_2]$ and S_{cyc} where S_{cyc} is the largest linear subset of S with no roots:

Lemma 1. (Store splitting)

Let S be a linear store, $\Gamma_1 \cap \Gamma_2 = \emptyset$ and $\Gamma_1, \Gamma_2 = \text{roots}(S)$. Then $S[\Gamma_1]$, $S[\Gamma_2]$ and S_{cyc} are pairwise disjoint linear stores and $S = S[\Gamma_1], S[\Gamma_2], S_{\text{cyc}}$.

Proof. We first show that they are pairwise disjoint: We use lemma 2 (see below) to prove that for any $x \in \text{dom}(S[\Gamma_1])$, we have $x \notin \text{dom}(S[\Gamma_2])$ and $x \notin \text{dom}(S_{\text{cyc}})$. By symmetry, the same holds with Γ_1 and Γ_2 swapped which implies pairwise disjointness.

Case Let $x \in \Gamma_1$. As x is a root of S , it does not occur in $\text{fv}(\text{rng}(S))$. As such, if $x \in \text{dom}(S[\Gamma_2])$ or $x \in \text{dom}(S_{\text{cyc}})$, it would be a root of $S[\Gamma_2]$ or S_{cyc} respectively. But $x \notin \Gamma_2$ by assumption and $x \notin \emptyset$. Thus the claim holds for all roots Γ_1 .

Case Let $x \in \text{fv}(v)$ with $y \mapsto v \in S[\Gamma_1]$, $y \notin \text{dom}(S[\Gamma_2])$ and $y \notin \text{dom}(S_{\text{cyc}})$. Since S is linear and $x \in \text{fv}(v)$, $x \notin \text{fv}(\text{rng}(S[\Gamma_2]))$ and $x \notin \text{fv}(\text{rng}(S_{\text{cyc}}))$. Then x might still be a root of $S[\Gamma_2]$ or S_{cyc} . But $x \notin \Gamma_2$ since x is not a root of S and $x \notin \emptyset$.

Now we show that they cover S . Let $S' = S - (S[\Gamma_1], S[\Gamma_2], S_{\text{cyc}})$. Then S' is a linear store with $\text{roots}(S') = \emptyset$:

Case S' is sound: Assume that $x \in \text{fv}(\text{rng}(S'))$ but $x \notin \text{dom}(S')$. Since S is linear, x must be a root of $S[\Gamma_1], S[\Gamma_2], S_{\text{cyc}}$. But by lemma 3 (see below), the roots of $S[\Gamma_1], S[\Gamma_2], S_{\text{cyc}}$ are Γ_1, Γ_2 which are also the roots of S . This is a contradiction, since x is not a root of S .

Case S' is linear: Clearly, since S' is a subset of S which is linear.

Case Assume that $x \in \text{roots}(S')$. Since $\Gamma_1, \Gamma_2 \in \text{dom}(S[\Gamma_1], S[\Gamma_2], S_{\text{cyc}})$, x is not a root of S . But since S is linear, this implies that $S[\Gamma_1], S[\Gamma_2], S_{\text{cyc}}$ is not sound, which is a contradiction to lemma 3. But then (by lemma 3, see below) we could add S' to S_{cyc} to obtain a linear subset of S with no roots. Since S_{cyc} is already the largest such subset, $S' = \emptyset$.

The above lemma makes use of the fact that $S[\Gamma]$ contains just values reachable from the roots Γ . We call this property *cycle-free*. Formally, a store S is cycle free if for any property P with

- $P(x)$ for all $x \in \Gamma$
- $P(x)$ for all $x \in \text{fv}(v)$ with $y \mapsto v \in S$ and $P(y)$

we have $P(x)$ for all $x \in \text{dom}(S)$. In particular, the induced store is cycle-free:

Lemma 2. (*Reachability from roots*)

Let S be a linear store, $\Gamma \subseteq \text{roots}(S)$ and $S[\Gamma]$ the smallest linear subset of S containing Γ . Then $S[\Gamma]$ is cycle-free.

Proof. Let S' be the subset of $S[\Gamma]$ for which P is not true. We claim that $S[\Gamma] - S'$ is still a linear subset of S containing Γ . Since $S[\Gamma]$ being the smallest such set, this implies $S' = \emptyset$.

Case Soundness of $S[\Gamma] - S'$: Assume that there is $x \in \text{fv}(\text{rng}(S[\Gamma] - S'))$ with $x \notin \text{dom}(S[\Gamma] - S')$. Choose $y \mapsto v \in (S[\Gamma] - S')$ with $x \in \text{fv}(v)$. Since $y \in \text{dom}(S[\Gamma] - S')$, we have $P(y)$. But since $S[\Gamma]$ is sound, $x \in \text{dom}(S[\Gamma])$ and thus $x \in \text{dom}(S')$. If $x \in \text{dom}(S')$, $P(x)$ is not true, which is a contradiction to the second rule of P .

Case Linearity of $S[\Gamma] - S'$: Since $S[\Gamma]$ is linear, so is any subset of $S[\Gamma]$.

Case $S[\Gamma] - S'$ contains Γ : By assumption $P(x)$ is true for all $x \in \Gamma$, so $\Gamma \cap S' = \emptyset$.

We can combine any linear stores again to obtain a linear store:

Lemma 3. (*Store joining*)

Let S_1 and S_2 be disjoint sound/linear stores with roots Γ_1 and Γ_2 . Then S_1, S_2 is a sound/linear store with roots Γ_1, Γ_2 .

Proof.

Case Soundness: Let $x \in \text{fv}(\text{rng}(S_1, S_2))$. Then either $x \in \text{fv}(\text{rng}(S_1))$ or $x \in \text{fv}(\text{rng}(S_2))$ and by the soundness of S_1 and S_2 , we have $x \in \text{dom}(S_1)$ or $x \in \text{dom}(S_2)$.

Case Linearity: Let $x \in \text{dom}(S_1, S_2)$. Since S_1 and S_2 are disjoint, either $x \in \text{dom}(S_1)$ (exclusive) or $x \in \text{dom}(S_2)$. Without loss of generality, assume $x \in \text{dom}(S_1)$. Then $x \notin \text{dom}(S_2)$ and by soundness of S_2 , $x \notin \text{fv}(\text{rng}(S_2))$. By linearity of S_1 , x occurs at most once in the free variables of

$\text{rng}(S_1)$. Since it does not occur in the free variables of $\text{rng}(S_2)$, it also occurs at most once in the free variables of $\text{rng}(S_1, S_2)$.

Case Roots: Since S_1 is sound and disjoint from S_2 , we have $\Gamma_2 \cap \text{fv}(\text{rng}(S_1)) = \emptyset$. By symmetry, $\Gamma_1, \Gamma_2 \subseteq \text{roots}(S_1, S_2)$. Assume $x \in \text{roots}(S_1, S_2)$. Then $x \in \text{roots}(S_1)$ or $x \in \text{roots}(S_2)$ and so $x \in \Gamma_1, \Gamma_2$.

B.2 Simple Invariant

Our simple invariant maintains that the store remains “well-formed” during evaluation. We use two stores: a “borrowed” store S_1 which is unchanged by evaluation and an “owned” store which can be changed. Our simple invariant $I(e, \Delta, \Gamma, S_1, S_2)$ is defined as:

- $\Delta \mid \Gamma \vdash e$
- S_1, S_2 are disjoint stores
- $\Delta \subseteq \text{dom}(S_1)$ and S_1 sound
- $\Gamma = \text{roots}(S_2)$ and S_2 linear

That invariant makes it safe to modify S_2 , as all values that the store semantics destroys are in Γ and not used anywhere else in the store ($\Gamma = \text{roots}(S_2)$). It would be enough for soundness to demand $\Gamma \subseteq \text{roots}(S_2)$. However, the stronger assertion directly gives us the garbage-free theorem and by the weakening lemma 4 we can always add separated memory to the store later:

Lemma 4. (*Weakening for store semantics*)

If $S \mid e \longrightarrow_s^* S' \mid e'$ then $S, S_1 \longrightarrow_s^* S', S_1 \mid e'$ for any S_1 with $\text{dom}(S_1) \cap \text{dom}(S) = \emptyset$.

Proof. By straight-forward induction on the judgement $S \mid e \longrightarrow_s^* S' \mid e'$.

B.3 Progress

In this section we want to show that the store semantics can progress if the simple invariant is true and operational semantics can progress. We assume throughout that for any function $f(\bar{y}; \bar{x}) = e \in \Sigma$, we have $\text{fv}(e) \subseteq \bar{y}, \bar{x}$. This is true if Σ is fully in-place:

Lemma 5. (*Free variables of FIP expressions are in Δ, Γ*)

If $\Delta \mid \Gamma \vdash e$, then $\text{fv}(e) \subseteq \Delta, \Gamma$.

Proof. By induction on the judgement $\Delta \mid \Gamma \vdash e$.

Case VAR:

- $\Delta \mid x \vdash x$ (1), given
- $\text{fv}(x) \subseteq \{x\}$ (2), definition

Case ATOM:

- $\Delta \mid \emptyset \vdash C$ (1), given
- $\text{fv}(C) = \emptyset$ (2), definition

Case TUPLE:

- $\Delta \mid \Gamma_1, \dots, \Gamma_n \vdash (v_1, \dots, v_n)$ (1), given
- $\Delta \mid \Gamma_i \vdash v_i$ (2), by TUPLE
- $\text{fv}(v_i) \subseteq \Delta, \Gamma_i$ (3), inductive hypothesis
- $\text{fv}((v_1, \dots, v_n)) = \text{fv}(v_1), \dots, \text{fv}(v_n) \subseteq \Delta, \Gamma_1, \dots, \Gamma_n$ (4), definition

Case REUSE:

$\Delta \mid \Gamma_1, \dots, Gk \vdash C^k v_1 \dots v_k$ (1), given
 $\Delta \mid \Gamma_i \vdash v_i$ (2), by REUSE
 $\text{fv}(v_i) \subseteq \Delta, \Gamma_i$ (3), inductive hypothesis
 $\text{fv}(C^k v_1 \dots v_k) = \text{fv}(v_1), \dots, \text{fv}(v_k) \subseteq \Delta, \Gamma_1, \dots, Gk$ (4), definition

Case CALL:

$\Delta \mid \Gamma \vdash f(\bar{y}; e)$ (1), given
 $\bar{y} \in \Delta, \text{dom}(\Sigma)$ (2), by CALL
 $\Delta \mid \Gamma \vdash e$ (3), by CALL
 $\text{fv}(e) \subseteq \Delta, \Gamma$ (4), inductive hypothesis
 $\text{fv}(f(\bar{y}; e)) = \bar{y}, \text{fv}(e) \subseteq \Delta, \Gamma$ (4), definition and $\text{dom}(\Sigma) \cap \text{fv}(e) = \emptyset$

Case BAPP:

$\Delta \mid \Gamma \vdash y e$ (1), given
 $y \in \Delta$ (2), by BAPP
 $\Delta \mid \Gamma \vdash e$ (3), by BAPP
 $\text{fv}(e) \subseteq \Delta, \Gamma$ (4), inductive hypothesis
 $\text{fv}(y e) = y, \text{fv}(e) \subseteq \Delta, \Gamma$ (4), definition

Case EMPTY:

$\Delta \mid \Gamma, \diamond_0 \vdash e$ (1), given
 $\Delta \mid \Gamma \vdash e$ (2), by EMPTY
 $\text{fv}(e) \subseteq \Delta, \Gamma$ (3), inductive hypothesis

Case LET:

$\Delta \mid \Gamma_1, \Gamma_2, \Gamma_3 \vdash \text{let } \bar{x} = e_1 \text{ in } e_2$ (1), given
 $\Delta, \Gamma_2 \mid \Gamma_1 \vdash e_1$ (2), by LET
 $\Delta \mid \Gamma_2, \Gamma_3, \bar{x} \vdash e_2$ (3), by LET
 $\text{fv}(e_1) \subseteq \Delta, \Gamma_2, \Gamma_1$ (4), inductive hypothesis
 $\text{fv}(e_2) \subseteq \Delta, \Gamma_2, \Gamma_3, \bar{x}$ (5), inductive hypothesis
 $\text{fv}(\text{let } \bar{x} = e_1 \text{ in } e_2) = \text{fv}(e_1), \text{fv}(e_2) - \bar{x} \subseteq \Delta, \Gamma_1, \Gamma_2, \Gamma_3$ (6), definition

Case BMATCH:

$\Delta \mid \Gamma \vdash \text{match } y \{ C_i \bar{x}_i \mapsto e_i \}$ (1), given
 $y \in \Delta$ (2), by BMATCH
 $\Delta, \bar{x}_i \mid \Gamma \vdash e_i$ (3), by BMATCH
 $\text{fv}(e_i) \subseteq \Delta, \bar{x}_i, \Gamma$ (4), inductive hypothesis
 $\text{fv}(\text{match } y \{ C_i \bar{x}_i \mapsto e_i \}) = y, \text{fv}(e_1) - \bar{x}_1, \dots, \text{fv}(e_n) - \bar{x}_n \subseteq \Delta, \Gamma$ (5), definition

Case DMATCH!:

$\Delta \mid \Gamma, x \vdash \text{match! } x \{ C_i \bar{x}_i \mapsto e_i \}$ (1), given
 $\Delta \mid \Gamma, \bar{x}_i, \diamond_k \vdash e_i$ (2), by DMATCH!
 $\text{fv}(e_i) \subseteq \Delta, \Gamma, \bar{x}_i$ (3), inductive hypothesis
 $\text{fv}(\text{match! } x \{ C_i \bar{x}_i \mapsto e_i \}) = x, \text{fv}(e_1) - \bar{x}_1, \dots, \text{fv}(e_n) - \bar{x}_n \subseteq \Delta, \Gamma$ (4), definition

We define $[S - \bar{x}]e$ as the substitution which replaces every variable $y \in \text{fv}(e) - \bar{x}$ by $[S]y$. Then:

Lemma 6. (*Store Substitution on unused variables*)

If $x \notin \text{fv}(e)$ or $x \notin \text{dom}(S)$, then $[S]e = [S - x]e$.

Proof. By induction on e .

Case Variable: $y \neq x$

$[S - x]y = [S]y$ (1), by definition

Case Variable: x

- $x \in \text{fv}(e)$ (1), since case necessary
- $x \notin \text{dom}(S)$ (2), by (1)
- $[S]x = x$ (3), by (2)
- $[S - x]x = x$ (4), by definition

Lemma 7. (*Store Substitution commutes*)

If S sound and $[S]\bar{z} = \bar{v}$, then $[S](e[\bar{x}:=\bar{z}]) = ([S - \bar{x}]e)[\bar{x}:=\bar{v}]$.

Proof.

- $[S]\bar{z} = \bar{v} = [S]\bar{v}$ (1), store substitution is idempotent
- $[S](e[\bar{x}:=\bar{z}]) = [S](e[\bar{x}:=\bar{v}])$ (2), by (1)
- $\text{fv}(\bar{v}) \cap \text{dom}(S) = \emptyset$ (3), definition
- $\text{fv}(\bar{v}) \subseteq \bar{z}$ (4), since S is sound
- $\bar{x} \cap \text{fv}(e[\bar{x}:=\bar{v}]) \subseteq \text{fv}(\bar{v})$ (5), since \bar{x} is substituted
- $[S](e[\bar{x}:=\bar{v}]) = [S - \bar{x}](e[\bar{x}:=\bar{v}])$ (6), lemma 6

Lemma 8. (*Store semantics reads values*)

If $I(\bar{v}, \Delta, \Gamma, S_1, S_2)$ then $S_1, S_2 \mid \bar{v} \longrightarrow_s^* S_1, S'_2 \mid \bar{x}$ with $[S_2]\bar{v} = [S'_2]\bar{x}$ and all names in $\text{dom}(S'_2) - \text{dom}(S_2)$ are fresh.

Proof. Using the $(x_1, \dots, \square, \dots, v_n)$ context, view each value v individually. By induction on v .

Case x : clear

- $S_1, S_2 \mid x \longrightarrow_s^* S_1, S_2 \mid x$ (1), reflexivity
- $x \in \text{dom}(S_2)$ (2), by invariant

Case C :

- $S_1, S_2 \mid C \longrightarrow_s S_1, S_2, x \mapsto C \mid x$ (1), (*atom_s*), fresh x
- $[S_2, x \mapsto C]x = [S_2]C$ (2), obvious

Case $C^k v_1 \dots v_k$:

- $\Delta \mid \Gamma_1, \dots, \Gamma_k, \diamond_k \vdash C^k v_1 \dots v_k$ (1), by REUSE
- $\Delta \mid \Gamma_1, \dots, \Gamma_k \vdash (v_1, \dots, v_k)$ (2), by TUPLE
- $I((v_1, \dots, v_k), \Delta, (\Gamma_1, \dots, \Gamma_k), S_1, S_2 - \diamond_k)$ (3), by (2)
- $S_1, (S_2 - \diamond_k) \mid (v_1, \dots, v_k) \longrightarrow_s^* S_1, S'_2 \mid \bar{x}$ (4), inductive hypothesis
- $[S'_2]\bar{x} = [S_2 - \diamond_k](v_1, \dots, v_k)$ (5), inductive hypothesis
- $S_1, S_2 \mid (v_1, \dots, v_k) \longrightarrow_s^* S_1, S'_2, \diamond_k \mid \bar{x}$ (6), weakening (4)
- $S_1, S'_2, \diamond_k \mid C^k \bar{x} \longrightarrow_s S_1, S'_2, x \mapsto C^k \bar{x} \mid x$ (7), (*reuse_s*), fresh x
- $[S'_2, x \mapsto C^k \bar{x}']x = [S_2](C^k v_1 \dots v_k)$ (8), by (5) and x fresh

We write $\text{drop } \bar{x}; e$ as a short-hand for $\text{drop } x_1; \dots \text{drop } x_n; e$.

Lemma 9. (*Dropping can progress*)

If $I((\text{drop } \bar{x}; e), \Delta, \Gamma, S_1, S_2)$ then $S_1, S_2 \mid \text{drop } \bar{x}; e \longrightarrow_s^* S_1, S'_2 \mid e$ with $[S_1, S_2]e = [S_1, S'_2]e$.

Proof. By induction on $|S_2|$.

Case $S_2 = \emptyset$:

- $\Delta \mid \emptyset \vdash \text{drop } \bar{x}; e$ (1), since $\text{roots}(S_2) = \emptyset$
- $\bar{x} = \emptyset$ (2), by DROP
- $S_1, S_2 \mid \text{drop } \emptyset; e \longrightarrow_s^* S_1, S_2 \mid e$ (3), by reflexivity

Case $S_2 \neq \emptyset$. Let $\bar{x} = x, \bar{x}'$.

$\Delta \mid \Gamma, x, \bar{x}' \vdash \text{drop } \bar{x}; e$	(1), by assumption
$x \notin \Gamma$	(2), by (1) and the multiplicity of x in Γ, \bar{x}
$x \notin \text{fv}(\text{drop } \bar{x}'; e)$	(3), by (2) and lemma 5
$S_2 = S'_2, x \mapsto C^k \bar{y}$	(4), by assumption
$[S_1, S_2]e = [S_1, S'_2]e$	(5), by (3)
$S_1, S'_2, x \mapsto C^k \bar{y} \mid \text{drop } x; \text{drop } \bar{x}'; e \longrightarrow_s S_1, S'_2 \mid \text{drop } (\bar{y}, \bar{x}'); e$	(6), by ($dcon_s$)
$ S'_2 + 1 = S_2 $	(7), by definition
$\Delta \mid \Gamma, \bar{y}, \bar{x}' \vdash \text{drop } (\bar{y}, \bar{x}'); e$	(8), by (1) and DROP
$\Gamma, x, \bar{x}' = \text{roots}(S_2)$	(9), by assumption
$\Gamma, \bar{y}, \bar{x}' = \text{roots}(S'_2)$	(10), by (4)
$I((\text{drop } (\bar{y}, \bar{x}'); e), \Delta, (\Gamma, \bar{y}, \bar{x}'), S_1, S_2)$	(11), by (8) and (10)
$S_1, S'_2 \mid \text{drop } (\bar{y}, \bar{x}'); e \longrightarrow_s^* S_1, S''_2 \mid e$	(12), by inductive hypothesis
$[S_1, S'_2]e = [S_1, S''_2]e$	(13), by inductive hypothesis
$[S_1, S_2]e = [S_1, S''_2]e$	(14), by (5) and (13)

Since free variables and static functions are separate syntactic categories, we define store substitution $[S]e$ to not act on functions. This assumption simplifies our proof, but does not actually change its semantics, since by lemma 36, we know that the store never contains functions anyway. Then we can take steps in the store semantics yielding the same value as the operational semantics:

Lemma 10. (*Store semantics can progress (no eval ctx)*)

If $I(e, \Delta, \Gamma, S_1, S_2)$ and $[S_1, S_2]e \longrightarrow e'$, then $S_1, S_2 \mid e \longrightarrow_s^* S_1, S'_2 \mid e''$ with $e' = [S_1, S'_2]e''$.

Proof. By case-analysis on $[S_1, S_2]e \longrightarrow e'$.

Case (*let*):

$[S_1, S_2](\text{let } \bar{x} = \bar{v} \text{ in } e') = (\text{let } \bar{x} = [S_1, S_2]\bar{v} \text{ in } [S_1, S_2 - \bar{x}]e')$	(1), definition
$\longrightarrow ([S_1, S_2 - \bar{x}]e')[\bar{x} := ([S_1, S_2]\bar{v})]$	(2), by (<i>let</i>)
$S_2 = S_2[\Gamma_1], S_2[\Gamma_2], S_2[\Gamma_3], S_{\text{cyc}}$	(3), by assumption
$I(\bar{v}, \emptyset, \Gamma_1, \emptyset, S_2[\Gamma_1])$	(4), since $\emptyset \mid \Gamma_1 \vdash \bar{v}$
$I(\bar{v}, \emptyset, \Gamma_1, (S_1, S_2[\Gamma_2], S_2[\Gamma_3], S_{\text{cyc}}), S_2[\Gamma_1])$	(5), weakening (4)
$S_1, S_2 \mid \bar{v} \longrightarrow_s^* S_1, S'_2 \mid \bar{z}$	(6), lemma 8
$S'_2 = S'_2[\Gamma_1], S_2[\Gamma_2], S_2[\Gamma_3], S_{\text{cyc}}$	(7), by (5)
$[S_1, S'_2]\bar{z} = [S_1, S_2]\bar{v}$	(8), lemma 8
$[S_1, S'_2 - \bar{x}]e' = [S_1, S_2 - \bar{x}]e'$	(9), since all new bindings are fresh
$S_1, S'_2 \mid \text{let } \bar{x} = \bar{z} \text{ in } e' \longrightarrow_s S_1, S'_2 \mid e'[\bar{x} := \bar{z}]$	(10), (let_s)
$[S_1, S'_2](e'[\bar{x} := \bar{z}]) = ([S_1, S'_2 - \bar{x}]e')[\bar{x} := ([S_1, S'_2]\bar{z})]$	(11), by lemma 7
$([S_1, S'_2 - \bar{x}]e')[\bar{x} := ([S_1, S'_2]\bar{z})] = ([S_1, S_2 - \bar{x}]e')[\bar{x} := ([S_1, S_2]\bar{v})]$	(12), by (8) and (9)

Case (*call*):

$$\begin{aligned}
 [S_1, S_2](f(\bar{y}'; \bar{v})) &= f([S_1, S_2]\bar{y}'; [S_1, S_2]\bar{v}) && (1), \text{ definition} \\
 &\longrightarrow e[\bar{y}:=([S_1, S_2]\bar{y}'), \bar{x}:=([S_1, S_2]\bar{v})] && (2), \text{ by } (call) \text{ with } f(\bar{y}; \bar{x}) = e \in \Sigma \\
 I((f(\bar{y}'; \bar{v})), \Delta, \Gamma_1, S_1, S_2) &&& (3), \text{ by assumption} \\
 I(\bar{v}, \Delta, \Gamma_1, S_1, S_2) &&& (4), \text{ by } CALL \\
 S_1, S_2 \mid \bar{v} &\longrightarrow_s^* S_1, S'_2 \mid \bar{v} && (5), \text{ lemma 8} \\
 [S_1, S'_2]\bar{z} &= [S_1, S_2]\bar{v} && (6), \text{ lemma 8} \\
 [S_1, S'_2]\bar{y}' &= [S_1, S_2]\bar{y}' && (7), \text{ since all new bindings are fresh} \\
 S_1, S'_2 \mid f(\bar{y}'; \bar{z}) &\longrightarrow_s S_1, S'_2 \mid e[\bar{y}:=\bar{y}', \bar{x}:=\bar{z}] && (8), \text{ by } (call_s) \text{ with } f(\bar{y}; \bar{x}) = e \in \Sigma \\
 [S_1, S'_2](e[\bar{y}:=\bar{y}', \bar{x}:=\bar{z}]) &&& \\
 &= [S_1, S'_2 - \bar{y}, \bar{x}]e[\bar{y}:=([S_1, S'_2]\bar{y}'), \bar{x}:=([S_1, S'_2]\bar{z})] && (9), \text{ lemma 7} \\
 &= e[\bar{y}:=([S_1, S'_2]\bar{y}'), \bar{x}:=([S_1, S'_2]\bar{z})] && (10), \text{ since } \text{fv}(e) \in \bar{y}, \bar{x}, \text{dom}(S) \\
 &= e[\bar{y}:=([S_1, S_2]\bar{y}'), \bar{x}:=([S_1, S_2]\bar{v})] && (11), \text{ by (6) and (7)}
 \end{aligned}$$

Case (app):

$$\begin{aligned}
 [S_1, S_2]((f) \bar{v}) &= (f) ([S_1, S_2]\bar{v}) && (1), \text{ by 36} \\
 &\longrightarrow e[\bar{x}:=([S_1, S_2]\bar{v})] && (2), \text{ by } (app) \text{ with } f(\bar{x}) = e \in \Sigma \\
 I((f) \bar{v}, \Delta, \Gamma_1, S_1, S_2) &&& (3), \text{ by assumption} \\
 I(\bar{v}, \Delta, \Gamma_1, S_1, S_2) &&& (4), \text{ by } BAPP \\
 S_1, S_2 \mid \bar{v} &\longrightarrow_s^* S_1, S'_2 \mid \bar{z} && (5), \text{ lemma 8} \\
 [S_1, S'_2]\bar{z} &= [S_1, S_2]\bar{v} && (6), \text{ lemma 8} \\
 S_1, S'_2 \mid (f) \bar{z} &\longrightarrow_s S_1, S'_2 \mid e[\bar{x}:=\bar{z}] && (7), \text{ by } (call_s) \text{ with } f(\bar{y}; \bar{x}) = e \in \Sigma \\
 [S_1, S'_2](e[\bar{x}:=\bar{z}]) &= [S_1, S'_2 - \bar{x}]e[\bar{x}:=([S_1, S'_2]\bar{z})] && (8), \text{ lemma 7} \\
 &= e[\bar{x}:=([S_1, S'_2]\bar{z})] && (9), \text{ since } \text{fv}(e) \in \bar{y}, \bar{x}, \text{dom}(S) \\
 &= e[\bar{x}:=([S_1, S_2]\bar{v})] && (10), \text{ by (6)}
 \end{aligned}$$

Case (match):

$$\begin{aligned}
 [S_1, S_2](\text{match } y \{ \overline{p \mapsto e} \}) &&& \\
 &= (\text{match } (C \bar{v}) \{ \overline{p \mapsto [S_1, S_2]e} \}) && (1), \text{ definition, where } [S_1, S_2]y = C^k \bar{v} \\
 &\longrightarrow ([S_1, S_2 - \bar{y}]e_i)[\bar{y}:=\bar{v}] && (2), \text{ by } (match) \\
 I((\text{match } y \{ \overline{p \mapsto e} \}), \Delta, \Gamma_1, S_1, S_2) &&& (3), \text{ by assumption} \\
 y \mapsto C^k \bar{z} \in S_1 &&& (4), \text{ by (3)} \\
 [S_1, S_2]y &= [S_1, S_2](C^k \bar{z}) = C^k \bar{v} && (5), \text{ since } [S_1, S_2]y = C^k \bar{v} \\
 S, y \mapsto C^k \bar{z} \mid \text{match } y \{ \overline{p \mapsto e} \} &\longrightarrow_s S, y \mapsto C^k \bar{z} \mid e_i[\bar{y}:=\bar{z}] && (6), \text{ by } (bmatch_s) \\
 [S_1, S_2](e_i[\bar{y}:=\bar{z}]) &= ([S_1, S_2 - \bar{y}]e_i)[\bar{y}:=\bar{v}] && (7), \text{ lemma 7}
 \end{aligned}$$

Case (match!):

$$\begin{aligned}
 [S_1, S_2](\text{match! } x \{ \overline{p \mapsto e} \}) &&& \\
 &= (\text{match! } (C \bar{v}) \{ \overline{p \mapsto [S_1, S_2]e} \}) && (1), \text{ definition, where } [S_1, S_2]x = C^k \bar{v} \\
 &\longrightarrow ([S_1, S_2 - \bar{x}]e_i)[\bar{x}:=\bar{v}] && (2), \text{ by } (match!) \\
 I((\text{match! } x \{ \overline{p \mapsto e} \}), \Delta, \Gamma_1, S_1, S_2) &&& (3), \text{ by assumption} \\
 x \mapsto C^k \bar{z} \in S_2 &&& (4), \text{ by (3)} \\
 [S_1, S_2]x &= [S_1, S_2](C^k \bar{z}) = C^k \bar{v} && (5), \text{ since } [S_1, S_2]x = C^k \bar{v} \\
 S, x \mapsto C^k \bar{z} \mid \text{match! } x \{ \overline{p \mapsto e} \} &\longrightarrow_s S, \diamond_k \mid e_i[\bar{x}:=\bar{z}] && (6), \text{ by } (dmatch_s) \\
 [S_1, S_2](e_i[\bar{x}:=\bar{z}]) &= ([S_1, S_2 - \bar{x}]e_i)[\bar{x}:=\bar{v}] && (7), \text{ lemma 7}
 \end{aligned}$$

Case (drop):

$$\begin{aligned}
 I((\text{drop } \bar{x}; e), \Delta, \Gamma, S_1, S_2) &&& (1), \text{ given} \\
 S_1, S_2 \mid \text{drop } \bar{x}; e &\longrightarrow_s^* S_1, S'_2 \mid e && (2), \text{ by lemma 9} \\
 [S_1, S_2]e &= [S_1, S'_2]e && (3), \text{ by lemma 9}
 \end{aligned}$$

Case (free):

$I((\text{free } k; e), \Delta, \Gamma, S_1, S_2)$	(1), given
$\diamond_k \in S_2$	(2), given
$S_1, S_2, \diamond_k \mid \text{free } k; e \longrightarrow_s^* S_1, S_2 \mid e$	(3), by (<i>free_s</i>)
$[S_1, S_2, \diamond_k]e = [S_1, S_2]e$	(4), obvious

B.4 Store semantics preserves linearity and roots

In this section, we wish to maintain the invariant $I(e, \Delta, \Gamma, S_1, S_2)$ which makes lemma 10 work.

Lemma 11. (*Variable substitution preserves FIP typing*)

Assuming that \bar{y} does not occur in e . (1) If $\Delta \mid \Gamma, \bar{x} \vdash e$, then $\Delta \mid \Gamma, \bar{y} \vdash e[\bar{x}:=\bar{y}]$. (2) If $\Delta, \bar{x} \mid \Gamma \vdash e$, then $\Delta, \bar{y} \mid \Gamma \vdash e[\bar{x}:=\bar{y}]$.

Proof. By induction on the FIP judgement for any such \bar{x}, \bar{y} .

Case VAR:

$\Delta \mid x \vdash x$	given
$\Delta \mid y \vdash x[x:=y]$	given

Case LET:

$\Delta \mid \Gamma_1, \Gamma_2, \Gamma_3, \bar{x}_1, \bar{x}_2, \bar{x}_3 \vdash \text{let } \bar{z} = e_1 \text{ in } e_2$	given
$\Delta \mid \Gamma_2, \Gamma_3, \bar{x}_2, \bar{x}_3, \bar{z} \vdash e_2$	given
$\Delta \mid \Gamma_2, \Gamma_3, \bar{y}_2, \bar{y}_3, \bar{z} \vdash e_2[\bar{x}_2:=\bar{y}_2, \bar{x}_3:=\bar{y}_3]$	inductive hypothesis (1)
$\Delta, \Gamma_2, \bar{x}_2 \mid \Gamma_1, \bar{x}_1 \vdash e_1$	given
$\Delta, \Gamma_2, \bar{x}_2 \mid \Gamma_1, \bar{y}_1 \vdash e_1[\bar{x}_1:=\bar{y}_1]$	inductive hypothesis (1)
$\Delta, \Gamma_2, \bar{y}_2 \mid \Gamma_1, \bar{y}_1 \vdash e_1[\bar{x}_1:=\bar{y}_1, \bar{x}_2:=\bar{y}_2]$	inductive hypothesis (2)
$\Delta \mid \Gamma_1, \Gamma_2, \Gamma_3, \bar{y}_1, \bar{y}_2, \bar{y}_3 \vdash (\text{let } \bar{z} = e_1 \text{ in } e_2)[\bar{x}:=\bar{y}]$	above

Other cases are clear.

Lemma 12. (*The delta environment can be weakened*)

If $\Delta \mid \Gamma \vdash e$, then $\Delta, x \mid \Gamma \vdash e$ for any $x \notin \Gamma$.

Proof. By straight-forward induction on the judgement $\Delta \mid \Gamma \vdash e$. We have to require $x \notin \Gamma$ as we have to ensure $\Delta, x \cap \Gamma = \emptyset$.

The next lemma is the main lemma of our soundness proof. It says that the individual steps of the store semantics take sound/linear stores to sound/linear stores. As usual, the judgement \longrightarrow_s does not include the eval rule.

Lemma 13. (*Store semantics preserves linearity and roots (no eval ctx)*)

If $I(e, \Delta, \Gamma, S_1, S_2)$ and $S_1, S_2 \mid e \longrightarrow_s S_1, S_3 \mid e'$, then $I(e', \Delta', \Gamma', S_1, S_3)$.

Proof. By case analysis on the rules of the store semantics.

Case (let_s).

$S_3 := S_2$	define
$\Delta \mid \Gamma, \bar{y} \vdash \text{let } \bar{x} = \bar{y} \text{ in } e$	given
$\Gamma, \bar{y} = \text{roots}(S_2) = \text{roots}(S_3)$	given
$\Delta \mid \Gamma, \bar{x} \vdash e$	by LET
$\Delta \mid \Gamma, \bar{y} \vdash e[\bar{x}:=\bar{y}]$	by lemma 11

Case (call_s).

FP²: Fully in-Place Functional Programming

$S_3 := S_2$ define
 $\Delta, \bar{y}' \mid \bar{x}' \vdash f(\bar{y}'; \bar{x}')$ given
 $\bar{x}' = \text{roots}(S_2) = \text{roots}(S_3)$ given
 $\bar{y}' \subseteq \text{dom}(S_1)$ given
 $\bar{y} \mid \bar{x} \vdash e$ by DEFFUN and $f(\bar{y}; \bar{x}) = e \in \Sigma$
 $\bar{y}' \mid \bar{x}' \vdash e[\bar{x}:=\bar{x}', \bar{y}:=\bar{y}']$ by lemma 11

Case (*anon_s*).

$S_3 := S_2$ (1), define
 $\Delta, f \mid \bar{x}' \vdash f \bar{x}'$ (2), given
 $\bar{x}' = \text{roots}(S_2) = \text{roots}(S_3)$ (3), by (2)
 $\emptyset \mid \bar{x} \vdash e$ (4), by DEFFUN and $f(\bar{x}) = e \in \Sigma$
 $\emptyset \mid \bar{x}' \vdash e[\bar{x}:=\bar{x}']$ (5), by lemma 11

Case (*reuse_s*).

$\Delta \mid \diamond_k, \bar{x} \vdash C^k \bar{x}$ (1), given
 $\diamond_k \in \text{roots}(S_2)$ (2), by (1)
 $S_3 := S_2 - \diamond_k, x \mapsto C^k \bar{x}$ (3), define, well-defined by (2)
 $\Delta \mid x \vdash x$ (4), by VAR
 $\bar{x}, \diamond_k = \text{roots}(S_2)$ (5), by (1)
 $x = \text{roots}(S_3)$ (6), by (3),(5) and since x is fresh
 S_3 linear (7), by (6)

Case (*atom_s*).

$\Delta \mid \emptyset \vdash C$ (1), given
 $S_2 = \emptyset$ (2), by (1)
 $S_3 := x \mapsto C$ (3), define
 $\Delta \mid x \vdash x$ (4), by VAR
 $x = \text{roots}(S_3)$ (5), by (3)
 S_3 linear (6), since x is fresh

Case (*bmatch_s*).

$\Delta \mid \Gamma \vdash \text{match } y \{ \overline{p \mapsto e} \}$ (1), given
 $y \mapsto C^k \bar{y} \in \text{dom}(S_1)$ (2), given
 $\bar{y} \subseteq \text{dom}(S_1)$ (3), since S_1 is sound
 $\Delta, \bar{y} \mid \Gamma \vdash e[\text{bv}(p):=\bar{y}]$ (4), by (1) and 11
 $S_3 := S_2$ (5), define

Case (*dmatch_s*).

$S_3 := S_2 - x, \diamond_k$ (1), define
 $\Delta \mid \Gamma \vdash \text{match } x \{ \overline{p \mapsto e} \}$ (2), given
 $x \mapsto C^k \bar{y} \in \text{roots}(S_2)$ (3), given
 S_3 sound (4), by (3)
 $\bar{y} \subseteq \text{roots}(S_3)$ (5), since S_2 is linear
 $\Delta \mid \Gamma, \bar{y}, \diamond_k \vdash e[\text{bv}(p):=\bar{y}]$ (6), by (2) and 11

Case (*dcon_s*).

- $S_3 := S_2 - x$ (1), define
- $\Delta \mid \Gamma, x \vdash \text{drop } x; e$ (2), given
- $x \mapsto C^k \bar{y} \in \text{roots}(S_2)$ (3), given
- S_3 sound (4), by (3)
- $\bar{y} \subseteq \text{roots}(S_3)$ (5), since S_2 is linear
- $\Delta \mid \Gamma, \bar{y} \vdash \text{drop } \bar{y}; e$ (6), by (2), DROP

Case (free_s).

- $\Delta \mid \Gamma, \diamond_k \vdash \text{free } k; e$ (1), given
- $\diamond_k \in \text{roots}(S_2)$ (2), given
- $S_3 := S_2 - \diamond_k$ (3), define, well-defined by (2)

B.5 Main Invariant

We now want to generalize lemma 13 to arbitrary evaluation contexts. In particular, we might start with $I(E[e], \Delta, \Gamma, S_1, S_2)$, take a step $S_1, S_2 \mid e \longrightarrow_s S_1, S_3 \mid e'$ and want to obtain $I(E[e'], \Delta', \Gamma', S_1, S_3)$. But unfortunately, this does not work. Let's see where this breaks down. We define:

- $E = \text{let } \bar{x} = \square \text{ in drop } \bar{x}; y$
- $e = \text{match } y \{ C \bar{y} \mapsto f(\bar{y};) \}$

Then:

- We start with $I(E[e], \emptyset, \{y\}, \emptyset, S_2)$
- We have $y \mid \emptyset \vdash e$, so $I(e, \{y\}, \emptyset, S_2, \emptyset)$
- We step $S_2 \mid \text{match } y \{ C \bar{y} \mapsto f(\bar{y};) \} \longrightarrow_s S_2 \mid f(\bar{y};)$
- Now we have $\bar{y} \mid \emptyset \vdash f(\bar{y};)$, so $I(e', \{\bar{y}\}, \emptyset, S_2, \emptyset)$

But then we do not have:

- $I(E[f(\bar{y};)], \emptyset, \{y\}, \emptyset, S_2)$, because the LET-rule only allows us to borrow y , not \bar{y}
- $I(E[f(\bar{y};)], \emptyset, \{y, \bar{y}\}, \emptyset, S_2)$, because \bar{y} are not used as owned.
- $I(E[f(\bar{y};)], \{\bar{y}\}, \{y\}, \emptyset, S_2)$, because clearly $\bar{y} \notin \emptyset$.

The problem seems to be that the Δ environment of the nested computation can change in such a way, that we can not put the environment back together in our invariant. Therefore, we need a weaker assumption which is maintained across nested evaluation steps. In this weaker invariant, we maintain the split of the evaluation context explicitly and allow a new Δ' environment for the contained expression. We define $I_E(e, \Delta, \Gamma, S_1, S_2)$ by induction on E :

Case $E = \square$:

$$I_{\square}(e, \Delta, \Gamma, S_1, S_2) := I(e, \Delta, \Gamma, S_1, S_2)$$

Case $E[E'] = C^k x_1 \dots E' \dots v_k \mid (x_1, \dots, E', \dots, v_n) \mid E' e \mid x E' \mid f(\bar{y}; E') \mid \text{match } E' \{ \overline{p \mapsto e} \}$:

$$I_{E[E']}(e, \Delta, (\Gamma_i, \Gamma), S_1, (S_2[\Gamma_i], S'_2)) := I_{E'}(e, \Delta', \Gamma_i, (S_1, S'_2), S_2[\Gamma_i]) \text{ and } I(E[\square], \Delta, \Gamma, S_1, S'_2)$$

Case $E[E'] = \text{let } \bar{x} = E' \text{ in } e$:

$$I_{E[E']}(e, \Delta, (\Gamma_1, \Gamma_2, \Gamma_3), S_1, (S_2[\Gamma_1], S'_2)) := \\ I_{E'}(e, \Delta', \Gamma_1, (S_1, S'_2), S_2[\Gamma_1]) \\ \text{and } I(E[\square], \Delta, (\Gamma_2, \Gamma_3), S_1, S'_2)$$

The big difference between this new invariant and the statement $I(E[e], \Delta, \Gamma, S_1, S_2)$, is that we allow an arbitrary Δ' environment in the recursive case (as long as Δ' continues to be a subset of S_1, S'_2). This simply weakens the previous invariant:

Lemma 14. (*Weakening the invariant*)

If $I(E[e], \Delta, \Gamma, S_1, S_2)$, then $I_E(e, \Delta, \Gamma, S_1, S_2)$.

By induction on E :

Case $E = \square$:

- $I(e, \Delta, \Gamma, S_1, S_2)$ (1), given
- $I_{\square}(e, \Delta, \Gamma, S_1, S_2)$ (2), by definition

Case $E[E'] = C^k x_1 \dots E' \dots v_k \mid (x_1, \dots, E', \dots, v_n) \mid E' e \mid x E' \mid f(\bar{y}; E') \mid \text{match } E' \{ \overline{p \mapsto e} \}$:

- $I(E[E'[e]], \Delta, \Gamma, S_1, S_2)$ (1), given
- $\Gamma = \Gamma_i, \Gamma'$ and $S_2 = S_2[\Gamma_i], S'_2$ (2), by the relevant rule
- $I(E'[e], \Delta, \Gamma_i, S_1, S_2[\Gamma_i])$ (3), by (2)
- $I(E[()], \Delta, \Gamma', S_1, S'_2)$ (4), by (2)
- $I_{E'}(e, \Delta, \Gamma_i, S_1, S_2[\Gamma_i])$ (5), inductive hypothesis on (3)
- $I_{E[E']} (e, \Delta, \Gamma, S_1, S_2)$ (6), by definition

Case $E[E'] = \text{let } \bar{x} = E' \text{ in } e$:

- $I(E[E'[e]], \Delta, \Gamma, S_1, S_2)$ (1), given
- $\Gamma = \Gamma_1, \Gamma_2, \Gamma_3$ and $S_2 = S_2[\Gamma_1], S'_2$ (2), by LET
- $I(E'[e], (\Delta, \Gamma_2), \Gamma_1, (S_1, S'_2), S_2[\Gamma_1])$ (3), by (2)
- $I(E[()], \Delta, (\Gamma_2, \Gamma_3), S_1, S'_2)$ (4), by (2)
- $I_{E'}(e, (\Delta, \Gamma_2), \Gamma_1, (S_1, S'_2), S_2[\Gamma_1])$ (5), inductive hypothesis on (3)
- $I_{E[E']} (e, \Delta, \Gamma, S_1, S_2)$ (6), by definition

Crucially, we can also weaken the simple variant inside the main invariant:

Lemma 15. (*Weakening the invariant*)

If $I_E(E_2[e], \Delta, \Gamma, S_1, S_2)$, then $I_{E[E_2]}(e, \Delta, \Gamma, S_1, S_2)$.

By induction on E :

Case $E = \square$:

- $I_{\square}(E_2[e], \Delta, \Gamma, S_1, S_2)$ (1), given
- $I(E_2[e], \Delta, \Gamma, S_1, S_2)$ (2), by definition
- $I_{E_2}(e, \Delta, \Gamma, S_1, S_2)$ (3), by lemma 14
- $I_{\square[E_2]}(e, \Delta, \Gamma, S_1, S_2)$ (4), by (3)

Case $E[E'] = C^k x_1 \dots E' \dots v_k \mid (x_1, \dots, E', \dots, v_n) \mid E' e \mid x E' \mid f(\bar{y}; E') \mid \text{match } E' \{ \overline{p \mapsto e} \}$:

- $I_{E[E']} (E_2[e], \Delta, (\Gamma_i, \Gamma), S_1, (S_2[\Gamma_i], S'_2))$ (1), given
- $I_{E'}(E_2[e], \Delta', \Gamma_i, S_1, S_2[\Gamma_i])$ (2), by definition
- $I(E[()], \Delta, \Gamma, S_1, S'_2)$ (3), by definition
- $I_{E'[E_2]}(e, \Delta', \Gamma_i, S_1, S_2[\Gamma_i])$ (4), by inductive hypothesis on (2)
- $I_{E[E'[E_2]]}(e, \Delta, (\Gamma_i, \Gamma), S_1, (S_2[\Gamma_i], S'_2))$ (5), by (3) and (4)

Case $E[E'] = \text{let } \bar{x} = E' \text{ in } e$:

- $I_{E[E']} (E_2[e], \Delta, (\Gamma_1, \Gamma_2, \Gamma_3), S_1, (S_2[\Gamma_1], S'_2))$ (1), given
- $I_{E'}(E_2[e], \Delta', \Gamma_1, (S_1, S'_2), S_2[\Gamma_1])$ (2), by definition
- $I(E[()], \Delta, (\Gamma_2, \Gamma_3), S_1, S'_2)$ (3), by definition
- $I_{E'[E_2]}(e, \Delta', \Gamma_1, (S_1, S'_2), S_2[\Gamma_1])$ (4), by inductive hypothesis on (2)
- $I_{E[E'[E_2]]}(e, \Delta, (\Gamma_1, \Gamma_2, \Gamma_3), S_1, (S_2[\Gamma_1], S'_2))$ (5), by (3) and (4)

Of course, we do not want to remain in the weaker invariant forever. But how can we get back, if the Δ' environment is wrong? The trick is that we have to wait until execution of the nested expression is finished. Then, any resulting value can be typed with an empty Δ' environment:

Lemma 16. (*Values do not borrow*)

If $\Delta \mid \Gamma \vdash \bar{v}$, then $\emptyset \mid \Gamma \vdash \bar{v}$.

Proof. By induction on $\Delta \mid \Gamma \vdash \bar{v}$:

Case x : clear

$\Delta \mid x \vdash x$ (1), given

$\emptyset \mid x \vdash x$ (2), by VAR

Case C :

$\Delta \mid \emptyset \vdash C$ (1), given

$\emptyset \mid \emptyset \vdash C$ (2), by ATOM

Case $C^k v_1 \dots v_k$:

$\Delta \mid \Gamma_1, \dots, \Gamma_k, \diamond_k \vdash C^k v_1 \dots v_k$ (1), by REUSE

$\Delta \mid \Gamma_i \vdash v_i$ (2), by (1)

$\emptyset \mid \Gamma_i \vdash v_i$ (3), inductive hypothesis

$\emptyset \mid \Gamma_1, \dots, \Gamma_k, \diamond_k \vdash C^k v_1 \dots v_k$ (4), by REUSE

Case (v_1, \dots, v_k) :

$\Delta \mid \Gamma_1, \dots, \Gamma_k \vdash (v_1, \dots, v_k)$ (1), by TUPLE

$\Delta \mid \Gamma_i \vdash v_i$ (2), by (1)

$\emptyset \mid \Gamma_i \vdash v_i$ (3), inductive hypothesis

$\emptyset \mid \Gamma_1, \dots, \Gamma_k \vdash (v_1, \dots, v_k)$ (4), by TUPLE

Lemma 17. (*Values do not borrow*)

If $I(\bar{v}, \Delta, \Gamma, S_1, S_2)$, then $I(\bar{v}, \emptyset, \Gamma, S'_1, S_2)$ for any S'_1 .

Proof.

$I(\bar{v}, \Delta, \Gamma, S_1, S_2)$ (1), given

$\Delta \mid \Gamma \vdash \bar{v}$ (2), by definition

S_1, S_2 are disjoint stores (3), by definition

$\Delta \subseteq \text{dom}(S_1)$ and S_1 sound (4), by definition

$\Gamma = \text{roots}(S_2)$ and S_2 linear (5), by definition

$\emptyset \mid \Gamma \vdash \bar{v}$ (6), by lemma 16

$I(\bar{v}, \emptyset, \Gamma, S'_1, S_2)$ (7), above

However, notice that we introduce a Δ' environment at every level of the evaluation context. That means, that we can only remove a single level of the evaluation context at a time. To model this, we call an evaluation context E flat if $E = E'[E'']$ implies that either E' or E'' is the hole. Then:

Lemma 18. (*Strengthening the invariant*)

If E is flat and $I_E(\bar{v}, \Delta, \Gamma, S_1, S_2)$, then $I(E[\bar{v}], \Delta, \Gamma, S_1, S_2)$.

Proof. By case analysis on E .

Case $E = \square$:

$I_{\square}(\bar{v}, \Delta, \Gamma, S_1, S_2)$ (1), given

$I(\bar{v}, \Delta, \Gamma, S_1, S_2)$ (2), by definition

Case $E = C^k x_1 \dots \square \dots v_k \mid (x_1, \dots, \square, \dots, v_n) \mid \square e \mid x \square \mid f(\bar{y}; \square) \mid \text{match } \square \{ \bar{p} \mapsto e \}$:

FP²: Fully in-Place Functional Programming

$$\begin{aligned}
I_E(\bar{v}, \Delta, (\Gamma_i, \Gamma), S_1, (S_2[\Gamma_i], S'_2)) & \quad (1), \text{ given} \\
I_{\square}(\bar{v}, \Delta', \Gamma_i, S_1, S_2[\Gamma_i]) & \quad (2), \text{ by definition} \\
I(E[()], \Delta, \Gamma, S_1, S'_2) & \quad (3), \text{ by definition} \\
I(\bar{v}, \Delta', \Gamma_i, S_1, S_2[\Gamma_i]) & \quad (4), \text{ by (2)} \\
I(\bar{v}, \emptyset, \Gamma_i, (S_1, S'_2), S_2[\Gamma_i]) & \quad (5), \text{ by lemma 17} \\
I(E[\bar{v}], \Delta, (\Gamma_i, \Gamma), S_1, (S_2[\Gamma_i], S'_2)) & \quad (6), \text{ by (3) and (5)}
\end{aligned}$$

Case $E = \text{let } \bar{x} = \square \text{ in } e$:

$$\begin{aligned}
I_E(\bar{v}, \Delta, (\Gamma_1, \Gamma_2, \Gamma_3), S_1, (S_2[\Gamma_1], S'_2)) & \quad (1), \text{ given} \\
I_{\square}(\bar{v}, \Delta', \Gamma_1, (S_1, S'_2), S_2[\Gamma_1]) & \quad (2), \text{ by definition} \\
I(E[()], \Delta, (\Gamma_2, \Gamma_3), S_1, S'_2) & \quad (3), \text{ by definition} \\
I(\bar{v}, \Delta', \Gamma_1, (S_1, S'_2), S_2[\Gamma_1]) & \quad (4), \text{ by (2)} \\
I(\bar{v}, \emptyset, \Gamma_1, (S_1, S'_2), S_2[\Gamma_1]) & \quad (5), \text{ by lemma 17} \\
I(E[\bar{v}], \Delta, (\Gamma_1, \Gamma_2, \Gamma_3), S_1, (S_2[\Gamma_1], S'_2)) & \quad (6), \text{ by (3) and (5)}
\end{aligned}$$

Again, we can also strengthen inside the main invariant:

Lemma 19. (*Strengthening the invariant*)

If E_2 is flat and $I_{E[E_2]}(\bar{v}, \Delta, \Gamma, S_1, S_2)$, then $I_E(E_2[\bar{v}], \Delta, \Gamma, S_1, S_2)$.

Proof. By induction on E .

Case $E = \square$:

$$\begin{aligned}
I_{\square[E_2]}(\bar{v}, \Delta, \Gamma, S_1, S_2) & \quad (1), \text{ given} \\
I_{E_2}(\bar{v}, \Delta, \Gamma, S_1, S_2) & \quad (2), \text{ by definition} \\
I(E_2[\bar{v}], \Delta, \Gamma, S_1, S_2) & \quad (3), \text{ by lemma 18} \\
I_{\square}(E_2[\bar{v}], \Delta, \Gamma, S_1, S_2) & \quad (4), \text{ by (3)}
\end{aligned}$$

Case $E[E'] = C^k x_1 \dots E' \dots v_k \mid (x_1, \dots, E', \dots, v_n) \mid E' e \mid x E' \mid f(\bar{y}; E') \mid \text{match } E' \{ \bar{p} \mapsto \bar{e} \}$:

$$\begin{aligned}
I_{E[E_2]}(\bar{v}, \Delta, (\Gamma_i, \Gamma), S_1, (S_2[\Gamma_i], S'_2)) & \quad (1), \text{ given} \\
I(E[()], \Delta, \Gamma, S_1, S'_2) & \quad (2), \text{ by definition} \\
I_{E'}(\bar{v}, \Delta', \Gamma_i, (S_1, S'_2), S_2[\Gamma_i]) & \quad (3), \text{ by definition} \\
I_{E'}(E_2[\bar{v}], \Delta', \Gamma_i, (S_1, S'_2), S_2[\Gamma_i]) & \quad (4), \text{ inductive hypothesis on (3)} \\
I_{E[E']}(E_2[\bar{v}], \Delta, (\Gamma_i, \Gamma), S_1, (S_2[\Gamma_i], S'_2)) & \quad (5), \text{ by (2) and (4)}
\end{aligned}$$

Case $E[E'] = \text{let } \bar{x} = E' \text{ in } e$:

$$\begin{aligned}
I_{E[E_2]}(\bar{v}, \Delta, (\Gamma_1, \Gamma_2, \Gamma_3), S_1, (S_2[\Gamma_1], S'_2)) & \quad (1), \text{ given} \\
I(E[()], \Delta, (\Gamma_2, \Gamma_3), S_1, S'_2) & \quad (2), \text{ by definition} \\
I_{E'}(\bar{v}, \Delta', \Gamma_1, (S_1, S'_2), S_2[\Gamma_1]) & \quad (3), \text{ by definition} \\
I_{E'}(E_2[\bar{v}], \Delta', \Gamma_1, (S_1, S'_2), S_2[\Gamma_1]) & \quad (4), \text{ by inductive hypothesis on (3)} \\
I_{E[E']}(E_2[\bar{v}], \Delta, (\Gamma_1, \Gamma_2, \Gamma_3), S_1, (S_2[\Gamma_1], S'_2)) & \quad (5), \text{ by (2) and (4)}
\end{aligned}$$

B.6 Soundness

Now we can show the main soundness theorem. First, we extend the progress and preservation proofs to handle the `STEP` and `EVAL` cases. If we evaluate e_1 under the context E_1 , we need to assume the invariant $I_{E_1}(e_1, \Delta, \Gamma, S_1, S_2)$. But how can we obtain the invariant for this precise E_1 ? The trick is that we do not have to know E_1 , instead we can just assume $I_{E_2}(e_2, \Delta, \Gamma, S_1, S_2)$ for $E_1[e_1] = E_2[e_2]$:

Lemma 20. (*Comparing evaluation contexts*)

Let $E_1[e_1] = E_2[e_2]$, then either:

- $E_1 = E'_1[E''_1]$ with $E'_1 = E_2$ and $e_2 = E''_1[e_1]$
- $E_2 = E'_2[E''_2]$ with $E'_2 = E_1$ and $e_1 = E''_2[e_2]$

Proof. By induction on the pair (E_1, E_2) :

Case $E_1 = \square$:

$$\begin{aligned} e_1 &= E_2[e_2] \quad (1), \text{ given} \\ E_2 &= \square[E_2] \quad (2), \text{ obvious} \end{aligned}$$

Case $E_2 = \square$:

$$\begin{aligned} e_2 &= E_1[e_1] \quad (1), \text{ given} \\ E_1 &= \square[E_1] \quad (2), \text{ obvious} \end{aligned}$$

Case Else: $E_1 = E'_1[E''_1]$, $E_2 = E'_2[E''_2]$, and E'_1, E'_2 are flat.

$$\begin{aligned} E'_1[E''_1[e_1]] &= E'_2[E''_2[e_2]] && (1), \text{ given} \\ E'_1 &= E'_2 && (2), \text{ by (1)} \\ E''_1 &= E_3[E_4] \text{ with } E_3 = E''_2 \text{ and } e_2 = E_4[e_1] && (3), \text{ inductive hypothesis, wlog case (1) holds} \\ E_1 &= E'_1[E''_1] = E'_1[E_3[E_4]] && (4), \text{ by (3)} \\ E'_1[E_3] &= E'_2[E''_2] && (5), \text{ by (2) and (3)} \end{aligned}$$

Lemma 21. (*Alignment of invariant*)

If $e_1 \longrightarrow e'_1$, $E_1[e_1] = E_2[e_2]$ and $I_{E_2}(e_2, \Delta, \Gamma, S_1, S_2)$, then $I_{E_1}(e_1, \Delta, \Gamma, S_1, S_2)$.

Proof. Use lemma 20 to compare E_1 and E_2 .

Case $E_1 = E'_1[E''_1]$ with $E'_1 = E_2$ and $e_2 = E''_1[e_1]$

$$\begin{aligned} I_{E_2}(e_2, \Delta, \Gamma, S_1, S_2) & \quad (1), \text{ given} \\ I_{E'_1}(E''_1[e_1], \Delta, \Gamma, S_1, S_2) & \quad (2), \text{ by assumption} \\ I_{E'_1[E''_1]}(e_1, \Delta, \Gamma, S_1, S_2) & \quad (3), \text{ lemma 15} \\ I_{E_1}(e_1, \Delta, \Gamma, S_1, S_2) & \quad (4), \text{ by (3)} \end{aligned}$$

Case $E_2 = E'_2[E''_2]$ with $E'_2 = E_1$ and $e_1 = E''_2[e_2]$

$$\begin{aligned} I_{E_2}(e_2, \Delta, \Gamma, S_1, S_2) & \quad (1), \text{ given} \\ I_{E'_2[E''_2]}(e_2, \Delta, \Gamma, S_1, S_2) & \quad (2), \text{ by assumption} \\ e_1 = E''_2[e_2] & \longrightarrow e'_1 \quad (3), \text{ given} \end{aligned}$$

Use case analysis on (3). In every case, either $E''_2 = \square$ or E''_2 is flat with $e_2 = \bar{v}$.

If $E''_2 = \square$, then $E_2 = E'_2 = E_1$ and the claim holds by (1).

If E''_2 is flat with $e_2 = \bar{v}$, then

$$\begin{aligned} I_{E'_2}(E''_2[\bar{v}], \Delta, \Gamma, S_1, S_2) & \quad (4), \text{ by (2) and lemma 19} \\ I_{E_1}(e_1, \Delta, \Gamma, S_1, S_2) & \quad (5), \text{ by (4)} \end{aligned}$$

Lemma 22. (*Store semantics can progress*)

If $I_E(e, \Delta, \Gamma, S_1, S_2)$ and $[S_1, S_2]e \longrightarrow e'$, then $S_1, S_2 \mid e \longrightarrow_s^* S'_1, S'_2 \mid e''$ with $e' = [S_1, S'_2]e''$.

Proof.

$$\begin{aligned} I_E(e, \Delta, \Gamma, S_1, S_2) & \quad (1), \text{ given} \\ I(e, \Delta', \Gamma', S'_1, S'_2), \text{ with } S_1, S_2 = S'_1, S'_2 \text{ and } S'_2 \subseteq S_2 & \quad (2), \text{ by (1)} \\ [S_1, S_2]e = [S'_1, S'_2]e & \quad (3), \text{ by (2)} \\ S'_1, S'_2 \mid e \longrightarrow_s^* S'_1, S'_3 \mid e'' & \quad (4), \text{ by lemma 10} \\ e' = [S'_1, S'_3]e'' & \quad (5), \text{ by lemma 10} \\ S_3 := S'_3, (S_2 - S'_2) & \quad (6), \text{ define} \\ S_1, S_2 \mid e \longrightarrow_s^* S_1, S_3 \mid e'' & \quad (7), \text{ by (4)} \\ e' = [S_1, S_3]e'' & \quad (8), \text{ by (5)} \end{aligned}$$

Lemma 23. (*Store semantics preserves linearity and roots*)

If $I_E(e, \Delta, \Gamma, S_1, S_2)$ and $S_1, S_2 \mid e \longrightarrow_s S_1, S_3 \mid e'$, then $I_E(e', \Delta', \Gamma', S_1, S_3)$.

Proof.

- $I_E(e, \Delta, (\Gamma, \Gamma'), S_1, (S_2, S'_2))$ (1), given
- $I(e, \Delta', \Gamma, (S_1, S'_2), S_2)$ (2), split (1)
- $I_E(\cdot, \Delta, \Gamma', S_1, S'_2)$ (3), split (1)
- $S_1, S'_2, S_2 \mid e \longrightarrow_s S_1, S'_2, S'_3 \mid e'$ (4), by assumption
- $I(e', \Delta'', \Gamma'', (S_1, S'_2), S'_3)$ (5), by lemma 13
- $S_3 := S'_3, S'_2$ (6), define
- $I_E(e', \Delta, (\Gamma'', \Gamma'), S_1, S_3)$ (7), merge (3) and (5)

Lemma 24. (*Soundness lemma*)

If $I_E(e, \Delta, \Gamma, S_1, S_2)$ and $[S_1, S_2](E[e]) \longmapsto^* \bar{v}$, then $S_1, S_2 \mid E[e] \longmapsto_s^* S_1, S_3 \mid \bar{x}$ with $I(\bar{x}, \emptyset, \bar{x}, \emptyset, S_3)$ and $[S_3]\bar{x} = \bar{v}$.

Proof. By induction on $[S_1, S_2](E[e]) \longmapsto^* \bar{v}$.

Case Reflexive case:

- $[S_1, S_2](E[e]) = \bar{v}$ (1), given
- $E = \square, e = \bar{w}, [S_1, S_2]\bar{w} = \bar{v}$ (2), by (1)
- $I_\square(\bar{w}, \Delta, \Gamma, S_1, S_2)$ (3), given
- $I(\bar{w}, \Delta, \Gamma, S_1, S_2)$ (4), by definition
- $S_1, S_2 \mid \bar{w} \longrightarrow_s^* S_1, S'_2 \mid \bar{x}$ (5), by lemma 8
- $\bar{v} = [S_1, S_2]\bar{w} = [S_1, S'_2]\bar{x}$ (6), by lemma 8
- $I(\bar{x}, \Delta, \Gamma, S_1, S'_2)$ (7), by lemma 13
- $I(\bar{x}, \emptyset, \Gamma, \emptyset, S'_2)$ (8), by lemma 17

Case Transitive case:

- $[S_1, S_2](E[e]) \longmapsto^* e'$ (1), given
- $e' \longmapsto^* \bar{v}$ (2), given
- $E'_1[e'_1] = [S_1, S_2](E[e]), e'_1 \longrightarrow e'_2, e' = E'_1[e'_2]$ (3), by STEP
- $[S_1, S_2]E_1 := E'_1, [S_1, S_2]e_1 := e'_1$ (4), define
- $E_1[e_1] = E[e]$ (5), by (4)
- $I_E(e, \Delta, \Gamma, S_1, S_2)$ (6), given
- $I_{E_1}(e_1, \Delta, \Gamma, S_1, S_2)$ (7), by lemma 21
- $S_1, S_2 \mid e_1 \longrightarrow_s^* S_1, S'_2 \mid e''$ (8), by lemma 22
- $e'_2 = [S_1, S'_2]e''$ (9), by lemma 22
- $I_{E_1}(e'', \Delta', \Gamma', S_1, S'_2)$ (10), by lemma 23
- $[S_1, S'_2](E_1[e'']) = ([S_1, S'_2]E_1)([S_1, S'_2]e'')$ (11), commute
- $= E'_1[e'_2]$ (12), by (4) and (9)
- $= e' \longmapsto^* \bar{v}$ (13), by (3) and (2)
- $S_1, S'_2 \mid E_1[e''] \longmapsto_s^* S_1, S_3 \mid \bar{x}$ (14), inductive hypothesis
- $I(\bar{x}, \emptyset, \bar{x}, \emptyset, S_3)$ and $[S_3]\bar{x} = \bar{v}$ (15), inductive hypothesis
- $S_1, S_2 \mid E_1[e_1] \longmapsto_s^* S_1, S'_2 \mid E_1[e'']$ (16), EVAL on (8)
- $S_1, S_2 \mid E[e] \longmapsto_s^* S_1, S_3 \mid \bar{x}$ (17), append (16) and (14)

Theorem 7. (*FIP programs are sound in store semantics*)

If $\Delta \mid \Gamma \vdash e$ and given disjoint stores S_1, S_2 with $\Delta \subseteq \text{dom}(S_1)$, S_1 sound, $\Gamma = \text{roots}(S_2)$ and S_2 linear, and $[S_1, S_2]e \longmapsto^* \bar{v}$, then $S_1, S_2 \mid e \longrightarrow_s^* S_1, S_3 \mid \bar{x}$ where $[S_3]\bar{x} = \bar{v}$, $\bar{x} = \text{roots}(S_3)$ and S_3 is linear.

Proof.

- $\Delta \mid \Gamma \vdash e$ (1), given
- S_1, S_2 disjoint stores (2), given
- $\Delta \subseteq \text{dom}(S_1)$ and S_1 sound (3), given
- $\Gamma = \text{roots}(S_2)$ and S_2 linear (4), given
- $I(e, \Delta, \Gamma, S_1, S_2)$ (5), by (1)-(4)
- $e = E[e']$ (6), for some E, e'
- $I_E(e', \Delta, \Gamma, S_1, S_2)$ (7), lemma 14
- $S_1, S_2 \mid E[e'] \xrightarrow{*}_s S_1, S_3 \mid \bar{x}$ (8), lemma 24
- $I(\bar{x}, \emptyset, \bar{x}, \emptyset, S_3)$ (9), lemma 24
- $[S_3]\bar{x} = \bar{v}$ (10), lemma 24
- $\bar{x} = \text{roots}(S_3)$ and S_3 linear (11), by (9)

C STACK BOUND

In this section we wish to prove the stack bound on FIP algorithms.

C.1 Within a function

First, we show a bound on the evaluation context when no functions are called by working with a version of the store semantics without the (*anon_s*) and (*call_s*) rules. In practice, we would not allocate a stack frame for these parts of the evaluation context.

Below, we define the depth $|e|$ of an expression and the depth $|E|$ of an evaluation context. We fix a signature Σ and denote by $|e_{\max}|$ the maximum depth of an expression bound in Σ .

$ x $	$:= 0$
$ C^k v_1 \dots v_k $	$:= 1 + \max\{ v_1 , \dots, v_k \}$
$ (v_1, \dots, v_n) $	$:= 1 + \max\{ v_1 , \dots, v_n \}$
$ e_1 e_2 $	$:= 1 + \max\{ e_1 , e_2 \}$
$ f(e_1; e_2) $	$:= 1 + \max\{ e_1 , e_2 \}$
$ \text{let } \bar{x} = e_1 \text{ in } e_2 $	$:= 1 + \max\{ e_1 , e_2 \}$
$ \text{match } e \{ \overline{p \mapsto e} \} $	$:= 1 + \max\{ e , e_1 , \dots, e_n \}$
$ \text{match! } e \{ \overline{p \mapsto e} \} $	$:= 1 + \max\{ e , e_1 , \dots, e_n \}$
$ \text{drop } \bar{x}; e $	$:= 1 + e $
$ \text{free } k; e $	$:= 1 + e $
$ \square $	$:= 0$
$ C^k x_1 \dots E' \dots v_k $	$:= 1 + E' $
$ (x_1, \dots, E', \dots, v_n) $	$:= 1 + E' $
$ E' e $	$:= 1 + E' $
$ x E' $	$:= 1 + E' $
$ f(\bar{y}; E') $	$:= 1 + E' $
$ \text{let } \bar{x} = E' \text{ in } e $	$:= 1 + E' $
$ \text{match! } E \{ \overline{p \mapsto e} \} $	$:= 1 + E' $

We also re-define the free variables fv to also collect all directly called functions

$$\text{fv}(f(e_1; e_2)) := \{f\}, \text{fv}(e_1), \text{fv}(e_2)$$

Lemma 25. (*Relating the depths of expressions and eval contexts*)

We have $|E[e]| \geq |E| + |e|$.

Proof. By straight-forward induction on E .

Lemma 26. (*Substitution of variables leaves depth unchanged*)

We have $|e| = |e[\bar{x}:=\bar{y}]|$.

Proof. By straight-forward induction on e .

Lemma 27. (*Substitution inside eval context*)

If $|e_1| \geq |e_2|$, then $|E[e_1]| \geq |E[e_2]|$.

Proof. By straight-forward induction on E .

Case \square : By assumption.

Case $C^k x_1 \dots E' \dots v_k$:

$$\begin{aligned}
 |C^k x_1 \dots E'[e_1] \dots v_k| &= 1 + \max\{|x_1|, \dots, |E'[e_1]|, \dots, |v_k|\} & (1), \text{ by definition} \\
 |C^k x_1 \dots E'[e_2] \dots v_k| &= 1 + \max\{|x_1|, \dots, |E'[e_2]|, \dots, |v_k|\} & (2), \text{ by definition} \\
 |E'[e_1]| &\geq |E'[e_2]| & (3), \text{ inductive hypothesis} \\
 |E[e_1]| &\geq |E[e_2]| & (4), \text{ by (1),(2),(3)}
 \end{aligned}$$

Case $(x_1 \dots E' \dots v_k)$:

$$\begin{aligned}
 |(x_1 \dots E'[e_1] \dots v_k)| &= 1 + \max\{|x_1|, \dots, |E'[e_1]|, \dots, |v_k|\} & (1), \text{ by definition} \\
 |(x_1 \dots E'[e_2] \dots v_k)| &= 1 + \max\{|x_1|, \dots, |E'[e_2]|, \dots, |v_k|\} & (2), \text{ by definition} \\
 |E'[e_1]| &\geq |E'[e_2]| & (3), \text{ inductive hypothesis} \\
 |E[e_1]| &\geq |E[e_2]| & (4), \text{ by (1),(2),(3)}
 \end{aligned}$$

Case $E' e$:

$$\begin{aligned}
 |E'[e_1] e| &= 1 + \max\{|E'[e_1]|, |e|\} & (1), \text{ by definition} \\
 |E'[e_2] e| &= 1 + \max\{|E'[e_2]|, |e|\} & (2), \text{ by definition} \\
 |E'[e_1]| &\geq |E'[e_2]| & (3), \text{ inductive hypothesis} \\
 |E[e_1]| &\geq |E[e_2]| & (4), \text{ by (1),(2),(3)}
 \end{aligned}$$

Case $x E'$:

$$\begin{aligned}
 |x E'[e_1]| &= 1 + \max\{|x|, |E'[e_1]|\} & (1), \text{ by definition} \\
 |x E'[e_2]| &= 1 + \max\{|x|, |E'[e_2]|\} & (2), \text{ by definition} \\
 |E'[e_1]| &\geq |E'[e_2]| & (3), \text{ inductive hypothesis} \\
 |E[e_1]| &\geq |E[e_2]| & (4), \text{ by (1),(2),(3)}
 \end{aligned}$$

Case $f(\bar{y}; E')$:

$$\begin{aligned}
 |f(\bar{y}; E'[e_1])| &= 1 + \max\{|\bar{y}|, |E'[e_1]|\} & (1), \text{ by definition} \\
 |f(\bar{y}; E'[e_2])| &= 1 + \max\{|\bar{y}|, |E'[e_2]|\} & (2), \text{ by definition} \\
 |E'[e_1]| &\geq |E'[e_2]| & (3), \text{ inductive hypothesis} \\
 |E[e_1]| &\geq |E[e_2]| & (4), \text{ by (1),(2),(3)}
 \end{aligned}$$

Case let $\bar{x} = E'$ in e :

$$\begin{aligned}
 |\text{let } \bar{x} = E'[e_1] \text{ in } e| &= 1 + \max\{|E'[e_1]|, |e|\} & (1), \text{ by definition} \\
 |\text{let } \bar{x} = E'[e_2] \text{ in } e| &= 1 + \max\{|E'[e_2]|, |e|\} & (2), \text{ by definition} \\
 |E'[e_1]| &\geq |E'[e_2]| & (3), \text{ inductive hypothesis} \\
 |E[e_1]| &\geq |E[e_2]| & (4), \text{ by (1),(2),(3)}
 \end{aligned}$$

Case match! $E \{ \bar{p} \mapsto \bar{e} \}$:

$$\begin{aligned}
 |\text{match! } E'[e_1] \{ \bar{p} \mapsto \bar{e} \}| &= 1 + \max\{|E'[e_1]|, |e_1|, \dots, |e_n|\} & (1), \text{ by definition} \\
 |\text{match! } E'[e_2] \{ \bar{p} \mapsto \bar{e} \}| &= 1 + \max\{|E'[e_2]|, |e_1|, \dots, |e_n|\} & (2), \text{ by definition} \\
 |E'[e_1]| &\geq |E'[e_2]| & (3), \text{ inductive hypothesis} \\
 |E[e_1]| &\geq |E[e_2]| & (4), \text{ by (1),(2),(3)}
 \end{aligned}$$

Lemma 28. (Store semantics without calls reduces expression depth)

If $S | e \longrightarrow_s S' | e'$ not using the $(anon_s), (call_s)$ rules, then $|e| \geq |e'|$.

Proof. By induction on the evaluation context E of the evaluation.

Case $E = \square$:

Case (let_s) :

$$\begin{aligned}
 |\text{let } \bar{x} = \bar{y} \text{ in } e| &= 1 + \max\{|\bar{y}|, |e|\} & \text{definition} \\
 &> |e| = |e[\bar{x} := \bar{y}]| & \text{using lemma 26}
 \end{aligned}$$

Case $(reuse_s)$:

FP²: Fully in-Place Functional Programming

$$|C^k x_1 \dots x_k| = 1 + \max\{|x_1|, \dots, |x_k|\} = 1 > 0 = |x| \quad \text{definition}$$

Case ($atom_s$):

$$|C| = 1 > 0 = |x| \quad \text{definition}$$

Case ($bmatch_s$), ($dmatch_s$):

$$\begin{aligned} |\text{match } y \{ \overline{p} \mapsto \overline{e} \}| &= 1 + \max\{0, |e_1|, \dots, |e_n|\} && \text{definition} \\ &> |e_i| = |e_i[\overline{x} := \overline{y}]| && \text{using lemma 26} \end{aligned}$$

Case ($dcon_s$):

$$|\text{drop } x; e| = 1 + |e| = |\text{drop } \overline{x}; e| \quad \text{definition}$$

Case ($free_s$):

$$|\text{free } k; e| = 1 + |e| > |e| \quad \text{definition}$$

Case $E \neq \square$:

$$\begin{aligned} E'[e_1] &\geq E'[e_2] && \text{by inductive hypothesis} \\ E[e_1] &\geq E[e_2] && \text{using lemma 27} \end{aligned}$$

Lemma 29. (*Inside a function, the stack bound is $|e|$*)

At any intermediate step $S \mid e[\overline{y} := \overline{y}', \overline{x} := \overline{x}'] \xrightarrow{*}_s S' \mid E[e']$ not using the ($anon_s$), ($call_s$) rules, we have $|E| + |e'| \leq |e|$.

Proof.

$$\begin{aligned} |e| &= |e[\overline{y} := \overline{y}', \overline{x} := \overline{x}']| && \text{lemma 26} \\ |e[\overline{y} := \overline{y}', \overline{x} := \overline{x}']| &\geq |E[e']| && \text{repeatedly apply lemma 28} \\ |E[e']| &\geq |E| + |e'| && \text{lemma 25} \end{aligned}$$

C.2 First-order

Next, we want to show the usual stack bound in a first order language, omitting just the ($anon_s$) rule this time. As in the paper, we work with a signature Σ where all functions f are defined so that their (mutually) recursive calls are in tailposition $\mathcal{T}[f]$. However, in a purely first-order sense this tail context is too strict. This is because the side-condition $f_i \notin \text{fv}(e_0)$ also prevents us from storing functions pointers in e_0 . However, in a first-order language, these function pointers do not matter: We only care about direct calls. We write $\widetilde{\mathcal{T}}[f]$ for the tail context where the side-condition $f_i \notin \text{fv}(e_0)$ is replaced by $f_i \notin \text{calls}(e_0)$ and $\text{calls}(e_0)$ contains all the functions directly called in e_0 . Clearly, if $e = \mathcal{T}[f]$, then $e = \widetilde{\mathcal{T}}[f]$.

Next, we want to show that the evaluation context does not grow under tail calls. Essentially, this is because the holes in the tail context are precisely the complement of the holes in the evaluation context [Bour et al. 2021].

Lemma 30. (*The eval context is complementary to the (first order) tail context*)

If $e = \widetilde{\mathcal{T}}[f]$, $f \in \overline{f}$ and $e = E[f(\overline{y}; \overline{x})]$, then $E = \square$.

If $e = \widetilde{\mathcal{T}}[f]$, $e = E[e']$ and $E \neq \square$, then $\overline{f} \cap \text{calls}(e') = \emptyset$.

Proof. The eval and tail context are defined as:

$$\begin{aligned} E ::= & \square \mid C^k x_1 \dots E \dots v_k \mid (x_1, \dots, E, \dots, v_n) \mid E e \mid x E \mid f(\overline{y}; E) \mid \text{let } \overline{x} = E \text{ in } e \\ & \mid \text{match! } E \{ \overline{p} \mapsto \overline{e} \} \end{aligned}$$

and

$$\begin{aligned} \widetilde{\mathcal{T}}[\bar{f}] ::= e_0 \mid f_i(e_0; e_0) \mid \text{let } \bar{x} = e_0 \text{ in } \mathcal{T}[\bar{f}] \mid \text{match } e_0 \{ \overline{p_i \mapsto \mathcal{T}_i[\bar{f}]} \} \mid \text{match! } e_0 \{ \overline{p_i \mapsto \mathcal{T}_i[\bar{f}]} \} \\ \mid \text{drop } \bar{x}; \mathcal{T}[\bar{f}] \mid \text{free } k; \mathcal{T}[\bar{f}] \end{aligned} \quad (\text{where } f_i \notin \text{calls}(e_0))$$

Unifying $\mathcal{T}[\bar{f}] = E[f(\bar{y}; \bar{x})]$ by matching on $\mathcal{T}[\bar{f}]$ yields $E = \square$ and $\mathcal{T}[\bar{f}] = f(\bar{y}; \bar{x})$.

Unifying $\mathcal{T}[\bar{f}] = E[e']$ by matching on $\mathcal{T}[\bar{f}]$ yields $f_i \notin \text{calls}(e')$ for all $f_i \in \bar{f}$.

Lemma 31. (*Substitution preserves first order tail contexts*)

If $e = \widetilde{\mathcal{T}}[\bar{f}]$ then $e[\bar{x}:=\bar{y}] = \widetilde{\mathcal{T}}[\bar{f}]$.

Proof. By straight-forward induction on $\widetilde{\mathcal{T}}[\bar{f}]$ as the substitution does not act on direct calls.

Lemma 32. (*Store semantics preserves first order tail contexts*)

If $S \mid e \longrightarrow_s S' \mid e'$ not using the $(\text{anon}_s), (\text{call}_s)$ rules and $e = \widetilde{\mathcal{T}}[\bar{f}]$ then $e' = \widetilde{\mathcal{T}}[\bar{f}]$.

Proof. We use case analysis on $\widetilde{\mathcal{T}}[\bar{f}]$. Let us first assume that $e = e_0$ where $f_i \notin \text{calls}(e_0)$. By induction on the evaluation context E of the evaluation.

Case (let_s):

$$\begin{aligned} S \mid \text{let } \bar{x} = \bar{y} \text{ in } e &\longrightarrow_s S \mid e[\bar{x}:=\bar{y}] && \text{definition} \\ \bar{f} \cap \text{calls}(e) &= \emptyset && \text{by assumption} \\ \bar{f} \cap \text{calls}(e[\bar{x}:=\bar{y}]) &= \emptyset && \text{by lemma 31} \end{aligned}$$

Case (reuse_s):

$$\begin{aligned} S, \diamond_k \mid C^k x_1 \dots x_k &\longrightarrow_s S, x \mapsto C^k x_1 \dots x_k \mid x && \text{definition, (fresh } x, k \geq 1) \\ \text{calls}(x) &= \emptyset && \text{by definition} \end{aligned}$$

Case (atom_s):

$$\begin{aligned} S \mid C &\longrightarrow_s S, x \mapsto C \mid x && \text{definition, (fresh } x) \\ \text{calls}(x) &= \emptyset && \text{by definition} \end{aligned}$$

Case (bmatch_s):

$$\begin{aligned} S, y \mapsto C^k \bar{y} \mid \text{match } y \{ \overline{p \rightarrow e} \} &\longrightarrow_s S, y \mapsto C^k \bar{y} \mid e_i[\bar{x}:=\bar{y}] && \text{definition, } (p_i = C^k \bar{y}) \\ \bar{f} \cap \text{calls}(e_i) &= \emptyset && \text{by assumption} \\ \bar{f} \cap \text{calls}(e_i[\bar{x}:=\bar{y}]) &= \emptyset && \text{by lemma 31} \end{aligned}$$

Case (dmatch_s):

$$\begin{aligned} S, x \mapsto C^k \bar{y} \mid \text{match! } x \{ \overline{p \rightarrow e} \} &\longrightarrow_s S, \diamond_k \mid e_i[\bar{x}:=\bar{y}] && \text{definition, } (p_i = C^k \bar{x}) \\ \bar{f} \cap \text{calls}(e_i) &= \emptyset && \text{by assumption} \\ \bar{f} \cap \text{calls}(e_i[\bar{x}:=\bar{y}]) &= \emptyset && \text{by lemma 31} \end{aligned}$$

Case (dcon_s):

$$\begin{aligned} S, x \mapsto C^k \bar{x} \mid \text{drop } x; e &\longrightarrow_s S \mid \text{drop } \bar{x}; e && \text{definition} \\ \bar{f} \cap \text{calls}(e) &= \emptyset && \text{by assumption} \\ \bar{f} \cap \text{calls}(\text{drop } \bar{x}; e) &= \emptyset && \text{by definition} \end{aligned}$$

Case (free_s):

$$\begin{aligned} S, \diamond_k \mid \text{free } k; e &\longrightarrow_s S \mid e && \text{definition} \\ \bar{f} \cap \text{calls}(e) &= \emptyset && \text{by assumption} \end{aligned}$$

Case (eval):

FP²: Fully in-Place Functional Programming

$$\begin{aligned}
S \mid E[e] &\longrightarrow_s S' \mid E[e'] && (1), \text{ given} \\
\bar{f} \cap \text{calls}(E[e]) &= \emptyset && (2), \text{ by assumption} \\
S \mid e &\longrightarrow_s S' \mid e' && (3), \text{ (eval) rule} \\
\bar{f} \cap \text{calls}(e') &= \emptyset && (4), \text{ inductive hypothesis} \\
\bar{f} \cap \text{calls}(E[e']) &= \emptyset && (5), \text{ by (2) and (4)}
\end{aligned}$$

We can now assume that $\tilde{\mathcal{T}}[\bar{f}] \neq e_0$. Then split on outermost layer of the evaluation context E . First we consider the case where $E \neq \square$.

$$\begin{aligned}
S \mid E[e] &\longrightarrow_s S' \mid E[e'] && (1), \text{ given} \\
\bar{f} \cap \text{calls}(e) &= \emptyset && (2), \text{ by lemma 30} \\
S \mid e &\longrightarrow_s S' \mid e' && (3), \text{ by (eval)} \\
\bar{f} \cap \text{calls}(e') &= \emptyset && (4), \text{ by the first case of this proof}
\end{aligned}$$

We need to show that $E[e']$ is again a tail-context. We split on the remaining cases of E that match a tail-context $\tilde{\mathcal{T}}[\bar{f}] \neq e_0$:

Case $E = f_i(\bar{y}; \square)$, $f_i \in \bar{f}$:

$$f_i(\bar{y}; e') = \tilde{\mathcal{T}}[\bar{f}] \quad (5), \text{ by (4)}$$

Case $E = \text{let } \bar{x} = \square \text{ in } e_2$:

$$\begin{aligned}
\text{let } \bar{x} = e \text{ in } e_2 &= \tilde{\mathcal{T}}[\bar{f}] && (5), \text{ given} \\
e_2 &= \tilde{\mathcal{T}}'[\bar{f}] && (6), \text{ by (5)} \\
\text{let } \bar{x} = e' \text{ in } e_2 &= \tilde{\mathcal{T}}[\bar{f}] && (7), \text{ by (4) and (6)}
\end{aligned}$$

Case $E = \text{match! } \square \{ \overline{p \mapsto e} \}$:

$$\begin{aligned}
\text{match! } e \{ \overline{p \mapsto e} \} &= \tilde{\mathcal{T}}[\bar{f}] && (5), \text{ given} \\
e_i &= \tilde{\mathcal{T}}_i[\bar{f}] && (6), \text{ by (5)} \\
\text{match! } e' \{ \overline{p \mapsto e} \} &= \tilde{\mathcal{T}}[\bar{f}] && (7), \text{ by (4) and (6)}
\end{aligned}$$

Lastly, assume that $\tilde{\mathcal{T}}[\bar{f}] \neq e_0$ and $E = \square$. Then split on the cases of $\tilde{\mathcal{T}}[\bar{f}]$.

Case $f_i(\bar{y}; \bar{x})$: Impossible, since only the call rule (which we omitted) can reduce in this case.

Case $\text{let } \bar{x} = \bar{y} \text{ in } \tilde{\mathcal{T}}[\bar{f}]$:

$$\begin{aligned}
S \mid \text{let } \bar{x} = \bar{y} \text{ in } e &\longrightarrow_s S \mid e[\bar{x}:=\bar{y}] && \text{definition} \\
e &= \tilde{\mathcal{T}}[\bar{f}] && \text{assumption} \\
e[\bar{x}:=\bar{y}] &= \tilde{\mathcal{T}}[\bar{f}] && \text{by lemma 31}
\end{aligned}$$

Case $\text{match } y \{ \overline{p_i \mapsto \tilde{\mathcal{T}}_i[\bar{f}]} \}$:

$$\begin{aligned}
S, y \mapsto C^k \bar{y} \mid \text{match } y \{ \overline{p \mapsto e} \} &\longrightarrow_s S, y \mapsto C^k \bar{y} \mid e_i[\bar{x}:=\bar{y}] && \text{definition, } (p_i = C^k \bar{y}) \\
e_i &= \tilde{\mathcal{T}}[\bar{f}] && \text{assumption} \\
e_i[\bar{x}:=\bar{y}] &= \tilde{\mathcal{T}}[\bar{f}] && \text{by lemma 31}
\end{aligned}$$

Case $\text{match! } x \{ \overline{p_i \mapsto \tilde{\mathcal{T}}_i[\bar{f}]} \}$:

$$\begin{aligned}
S, x \mapsto C^k \bar{y} \mid \text{match! } x \{ \overline{p \mapsto e} \} &\longrightarrow_s S, \diamond_k \mid e_i[\bar{x}:=\bar{y}] && \text{definition, } (p_i = C^k \bar{x}) \\
e_i &= \tilde{\mathcal{T}}[\bar{f}] && \text{assumption} \\
e_i[\bar{x}:=\bar{y}] &= \tilde{\mathcal{T}}[\bar{f}] && \text{by lemma 31}
\end{aligned}$$

Case $\text{drop } x$; $\tilde{\mathcal{T}}[\bar{f}]$:

$$\begin{array}{ll} S, x \mapsto C^k \bar{x} \mid \text{drop } x; e \longrightarrow_s S \mid \text{drop } \bar{x}; e & \text{definition} \\ e = \tilde{\mathcal{T}}[\bar{f}] & \text{assumption} \\ \text{drop } \bar{x}; e = \tilde{\mathcal{T}}[\bar{f}] & \text{by definition} \end{array}$$

Case free k ; $\tilde{\mathcal{T}}[\bar{f}]$:

$$\begin{array}{ll} S, \diamond_k \mid \text{free } k; e \longrightarrow_s S \mid e & \text{definition} \\ e = \tilde{\mathcal{T}}[\bar{f}] & \text{assumption} \end{array}$$

The store semantics also does not insert calls outside the signature into e' . We show this by showing that it preserves some FIP typing regarding a signature Σ .

Lemma 33. (*Store semantics preserves FIP typing*)

If $\Delta \mid \Gamma \vdash_\Sigma e$ and $S \mid e \longrightarrow_s S' \mid e'$ then $\Delta' \mid \Gamma' \vdash_\Sigma e'$.

Proof. By induction on the evaluation context E in $e = E[e_1]$. Case \square : By case analysis on the rules of the store semantics.

Case (let_s).

$$\begin{array}{ll} \Delta \mid \Gamma, \bar{y} \vdash \text{let } \bar{x} = \bar{y} \text{ in } e & \text{given} \\ \Delta \mid \Gamma, \bar{x} \vdash e & \text{by LET} \\ \Delta \mid \Gamma, \bar{y} \vdash e[\bar{x} := \bar{y}] & \text{by lemma 11} \end{array}$$

Case (call_s), (anon_s).

$$\begin{array}{ll} \bar{y} \mid \bar{x} \vdash e & \text{by DEFFUN and } f(\bar{y}; \bar{x}) = e \in \Sigma \\ \bar{y}' \mid \bar{x}' \vdash e[\bar{x} := \bar{x}', \bar{y} := \bar{y}'] & \text{by lemma 11} \end{array}$$

Case (reuse_s), (atom_s).

$$\Delta \mid x \vdash x \quad \text{by VAR}$$

Case (bmatch_s).

$$\begin{array}{ll} \Delta \mid \Gamma \vdash \text{match } y \{ \overline{p \mapsto e} \} & \text{given} \\ \Delta, \text{bv}(p) \mid \Gamma \vdash e & \text{by MATCH} \\ \Delta, \bar{y} \mid \Gamma \vdash e[\text{bv}(p) := \bar{y}] & \text{by 11} \end{array}$$

Case (dmatch_n).

$$\begin{array}{ll} \Delta \mid \Gamma \vdash \text{match } x \{ \overline{p \mapsto e} \} & \text{given} \\ \Delta \mid \Gamma, \text{bv}(p), \diamond_k \vdash e & \text{by MATCH!} \\ \Delta \mid \Gamma, \bar{y}, \diamond_k \vdash e[\text{bv}(p) := \bar{y}] & \text{by 11} \end{array}$$

Case (dcon_n).

$$\begin{array}{ll} \Delta \mid \Gamma, x \vdash \text{drop } x; e & \text{given} \\ \Delta \mid \Gamma, \bar{y} \vdash \text{drop } \bar{y}; e & \text{by DROP} \end{array}$$

Case (free_n).

$$\begin{array}{ll} \Delta \mid \Gamma, \diamond_k \vdash \text{free } k; e & \text{given} \\ \Delta \mid \Gamma \vdash e & \text{by FREE} \end{array}$$

Case $C^k x_1 \dots E' \dots v_k \mid (x_1, \dots, E', \dots, v_n) \mid E' e \mid x E' \mid f(\bar{y}; E') \mid \text{match! } E' \{ \overline{p \mapsto e} \}$:

$$\begin{array}{ll} \Delta \mid \Gamma \vdash E[e] & (1), \text{ given} \\ \Delta \mid \Gamma_1 \vdash E'[e] & (2), \text{ by (1), where } \Gamma_1 \subseteq \Gamma \\ \Delta' \mid \Gamma'_1 \vdash E'[e'] & (3), \text{ inductive hypothesis} \\ \Delta, \Delta' \mid \Gamma - \Gamma_1, \Gamma'_1 \vdash E[e'] & (4), \text{ by (3)} \end{array}$$

Case let $\bar{x} = E'$ in e : (notice that we may put “children” of Γ_2 into Δ')

$$\begin{aligned} \Delta \mid \Gamma \vdash E[e] & \quad (1), \text{ given} \\ \Delta, \Gamma_2 \mid \Gamma_1 \vdash E'[e] & \quad (2), \text{ by (1), where } \Gamma_1, \Gamma_2 \subseteq \Gamma \\ \Delta', \Gamma_2 \mid \Gamma'_1 \vdash E'[e'] & \quad (3), \text{ inductive hypothesis with } \Delta' \cap \Gamma_2 = \emptyset \\ \Delta, \Delta' \mid \Gamma - \Gamma_1, \Gamma'_1 \vdash E[e'] & \quad (4), \text{ by (3)} \end{aligned}$$

We can now prove the main lemma of this section:

Lemma 34. (*First order stack bound*)

Let Σ be non-empty and fully-in-place such that for all functions f in Σ mutually recursive with \bar{f} we have $f(\bar{y}; \bar{x}) = \tilde{\mathcal{T}}[\bar{f}]$. If $e = \tilde{\mathcal{T}}[\bar{f}]$, $\Delta \mid \Gamma \vdash_\Sigma e$ and $|e| \leq |e_{\max}|$, then at any intermediate step $S \mid e \xrightarrow{*}_s S' \mid E'[e']$ not using the (*anon_s*) rule, we have $|E'| \leq |e_{\max}| \cdot |\Sigma|$.

Proof. By induction on Σ . In both the base and the inductive case, we use another induction on all subderivations of $S \mid e \xrightarrow{*}_s S' \mid E'[e']$. Our induction hypothesis is then:

- Either $\Sigma = \bar{f}$ or the lemma holds for all derivations on $\Sigma' \subseteq \Sigma$.
- The lemma holds for Σ and all derivations $S_2 \mid e_2 \xrightarrow{*}_s S' \mid E'[e']$ such that $S \mid e \xrightarrow{*}_s S_2 \mid e_2$, $e_2 = \tilde{\mathcal{T}}[\bar{f}]$, $\Delta' \mid \Gamma' \vdash_\Sigma e_2$ and $|e_2| \leq |e_{\max}|$.

Case Base case $S \mid e \xrightarrow{*}_s S \mid E'[e']$ with $e = E'[e']$: Clear, since $|E'| \leq |e| \leq |e_{\max}|$ and $|\Sigma| \geq 1$.

If $S \mid e \xrightarrow{*}_s S' \mid E'[e']$ even without the call rule, we have

$$\begin{aligned} |E'| & \leq |E'| + |e'| \quad \text{clear} \\ & \leq |E'[e']| \quad \text{by lemma 25} \\ & \leq |e| \quad \text{repeatedly apply lemma 28} \\ & \leq |e_{\max}| \quad \text{by assumption} \end{aligned}$$

Else, we choose the longest subderivation from $E[e]$ that does not involve the call rule $S \mid e \xrightarrow{*}_s S_g \mid E_g[g(\bar{y}_2; \bar{x}_2)]$. We consider two cases: Either g is in the recursive group \bar{f} , or it appears in Σ before \bar{f} .

Case g is in the recursive group, $e = \tilde{\mathcal{T}}[g]$:

$$\begin{aligned} E_g[g(\bar{y}_2; \bar{x}_2)] & = \tilde{\mathcal{T}}[g] & \text{by lemma 32} \\ E_g & = \square & \text{by lemma 30} \\ S_g \mid g(\bar{y}_2; \bar{x}_2) & \xrightarrow{s} S_g \mid e_g[\bar{y}:=\bar{y}_2; \bar{x}:=\bar{x}_2] & (\text{call}_s) \\ e_g & = \tilde{\mathcal{T}}[\bar{f}] & \text{assumption} \\ \bar{y} \mid \bar{x} \vdash e_g & & \text{since } \Sigma \text{ is fully-in-place} \\ |e_g| & \leq |e_{\max}| & \text{by definition of } e_{\max} \\ |E'| & \leq |e_{\max}| \cdot |\Sigma| & \text{by the inductive hypothesis (2)} \end{aligned}$$

Case g is not in the same recursive group, $g \in \Sigma - \bar{f}$:

If evaluation never leaves g and $E' = E_g[E'']$, then:

$$\begin{aligned} |E_g| & \leq |e_{\max}| & \text{since the derivation involves no } (\text{call}_s) \\ g(\bar{y}_2; \bar{x}_2) & = \tilde{\mathcal{T}}[g_s] & \text{definition of } \tilde{\mathcal{T}}[g_s], \text{ where } g_s \text{ is } g\text{'s recursive group} \\ \bar{y}_2 \mid \bar{x}_2 \vdash g(\bar{y}_2; \bar{x}_2) & & \text{CALL} \\ |g(\bar{y}_2; \bar{x}_2)| & = 1 \leq |e_{\max}| & \text{definition of length} \\ S_g \mid E_g[g(\bar{y}_2; \bar{x}_2)] & \xrightarrow{*}_s S' \mid E_g[E''] & \text{since we chose a subderivation} \\ S_g \mid g(\bar{y}_2; \bar{x}_2) & \xrightarrow{*}_s S' \mid E''[e'] & \text{by the } (\text{eval}) \text{ rule} \\ |E''| & \leq |e_{\max}| \cdot |\Sigma - \bar{f}| & \text{by the inductive hypothesis (1)} \\ |\Sigma - \bar{f}| + 1 & \leq |\Sigma| & \text{definition of } \Sigma - \bar{f} \\ |E'| & = |E_g[E'']| = |E_g| + |E''| \leq |e_{\max}| \cdot |\Sigma| & \text{as above} \end{aligned}$$

If evaluation leaves g and $S_g \mid E_g[g(\bar{y}_2; \bar{x}_2)] \xrightarrow{s} S'_g \mid E_g[\bar{x}_3]$, then:

$ E_g[g(ys_2; xs_2)] \leq e_{\max} $	since the derivation involves no ($call_s$)
$ E_g[xs_3] \leq e_{\max} $	substitution of variables does not increase length
$E_g[g(ys_2; xs_2)] = \tilde{\mathcal{T}}[\bar{f}]$	by lemma 32
$E_g[xs_3] = \tilde{\mathcal{T}}[\bar{f}]$	substitution of variables does not introduce calls
$\Delta' \mid \Gamma' \vdash E_g[xs_3]$	repeatedly apply lemma 33
$ E' \leq e_{\max} \cdot \Sigma $	by the inductive hypothesis (2)

Together, these results imply the well-known stack bound for first-order programs:

Lemma 35. (*The first-order stack bound is $|e_{\max}||\Sigma|$*)

Let Σ be fully in-place such that for all functions f in Σ mutually recursive with \bar{f} we have $f(\bar{y}; \bar{x}) = \tilde{\mathcal{T}}[\bar{f}]$. Then at any intermediate step $S \mid f(\bar{y}; \bar{x}) \longrightarrow_s^* S' \mid E[e']$ not using the ($anon_s$) rule, we have $|E| \leq |e_{\max}| \cdot |\Sigma|$.

Proof. Apply lemma 34 to $f(\bar{y}; \bar{x})$.

Σ nonempty	since $f(\bar{y}; \bar{x}) = e \in \Sigma$
$f(\bar{y}; \bar{x}) = \tilde{\mathcal{T}}[\bar{f}]$	definition of $\tilde{\mathcal{T}}[\bar{f}]$
$\bar{y} \mid \bar{x} \vdash f(\bar{y}; \bar{x})$	CALL
$ f(\bar{y}; \bar{x}) = 1 \leq e_{\max} $	definition of length

C.3 Second-order

Now, we want to get to the main novelty of the proof: Giving a stack bound for anonymous function calls. The proof idea is that our calculus is second-order: An anonymously called function can only make direct calls to other functions (which are mutually recursive, or defined before it). The reason for this is that an anonymous function only receives the store and owned arguments in Γ , but we make sure that functions are never in the store or the Γ environment. Then the function has no way to access a function pointer f defined before itself to call anonymously.

Lemma 36. (*The store and Γ never contain functions*)

Let $\Delta \mid \Gamma \vdash e, g \notin \Gamma, \text{rng}(S)$ and $S \mid e \longrightarrow_s S' \mid e'$. Then $\Delta' \mid \Gamma' \vdash e'$ and $g \notin \Gamma', \text{rng}(S')$.

Proof. By induction on the evaluation context E in $e = E[e_1]$. Case \square : By case analysis on the rules of the store semantics.

Case (let_s).

$\Delta \mid \Gamma, \bar{y} \vdash \text{let } \bar{x} = \bar{y} \text{ in } e$	given
$\Delta \mid \Gamma, \bar{x} \vdash e$	by LET
$\Delta \mid \Gamma, \bar{y} \vdash e[\bar{x} := \bar{y}]$	by lemma 11
$g \notin \Gamma, \bar{y}, \text{rng}(S)$	by assumption

Case ($call_s$), ($anon_s$).

$\bar{y}' \mid \bar{x}' \vdash f(\bar{y}'; \bar{x}')$	given
$\bar{y} \mid \bar{x} \vdash e$	by DEFFUN and $f(\bar{y}; \bar{x}) = e \in \Sigma$
$\bar{y}' \mid \bar{x}' \vdash e[\bar{x} := \bar{x}', \bar{y} := \bar{y}']$	by lemma 11
$g \notin \bar{x}', \text{rng}(S)$	by assumption

Case ($reuse_s$), ($atom_s$).

$\Delta \mid \bar{x} \vdash C^k x_1 \dots x_k$	by REUSE
$g \notin \bar{x}, \text{rng}(S)$	given
$\Delta \mid x \vdash x$	by VAR
$g \notin x, \text{rng}(S, x \mapsto C^k x_1 \dots x_k)$	since x fresh

Case ($bmatch_s$).

$\Delta \mid \Gamma \vdash \text{match } y \{ \overline{p \mapsto e} \}$ given
 $\Delta, \text{bv}(p) \mid \Gamma \vdash e$ by MATCH
 $\Delta, \bar{y} \mid \Gamma \vdash e[\text{bv}(p):=\bar{y}]$ by 11
 $g \notin \Gamma, \text{rng}(S)$ by assumption

Case ($dmatch_n$).

$\Delta \mid \Gamma \vdash \text{match } x \{ \overline{p \mapsto e} \}$ given
 $g \notin \Gamma, \text{rng}(S, x \mapsto C^k y_1 \dots y_k)$ given
 $\Delta \mid \Gamma, \text{bv}(p), \diamond_k \vdash e$ by MATCH!
 $\Delta \mid \Gamma, \bar{y}, \diamond_k \vdash e[\text{bv}(p):=\bar{y}]$ by 11
 $g \notin \Gamma, \bar{y}, \diamond_k, \text{rng}(S, \diamond_k)$ as above

Case ($dcon_n$).

$\Delta \mid \Gamma, x \vdash \text{drop } x; e$ given
 $g \notin \Gamma, x, \text{rng}(S, x \mapsto C^k y_1 \dots y_k)$ given
 $\Delta \mid \Gamma, \bar{y} \vdash \text{drop } \bar{y}; e$ by DROP
 $g \notin \Gamma, \bar{y}, \text{rng}(S)$ as above

Case ($free_n$).

$\Delta \mid \Gamma, \diamond_k \vdash \text{free } k; e$ given
 $g \notin \Gamma, \diamond_k, \text{rng}(S, \diamond_k)$ given
 $\Delta \mid \Gamma \vdash e$ by FREE
 $g \notin \Gamma, \text{rng}(S)$ as above

Case $C^k x_1 \dots E' \dots v_k \mid (x_1, \dots, E', \dots, v_n) \mid E' e \mid x E' \mid f(\bar{y}; E') \mid \text{match! } E' \{ \overline{p \mapsto e} \}$:

$\Delta \mid \Gamma \vdash E[e]$ (1), given
 $g \notin \Gamma, \text{rng}(S)$ (2), given
 $\Delta \mid \Gamma_1 \vdash E'[e]$ (3), by (1), where $\Gamma_1 \subseteq \Gamma$
 $\Delta' \mid \Gamma'_1 \vdash E'[e']$ (4), inductive hypothesis
 $g \notin \Gamma'_1, \text{rng}(S')$ (5), inductive hypothesis
 $\Delta, \Delta' \mid \Gamma - \Gamma_1, \Gamma'_1 \vdash E[e']$ (6), by (5)
 $g \notin \Gamma - \Gamma_1, \Gamma'_1, \text{rng}(S')$ (7), by (2),(5)

Case let $\bar{x} = E'$ in e : (notice that we may put “children” of Γ_2 into Δ')

$\Delta \mid \Gamma \vdash E[e]$ (1), given
 $g \notin \Gamma, \text{rng}(S)$ (2), given
 $\Delta, \Gamma_2 \mid \Gamma_1 \vdash E'[e]$ (3), by (1), where $\Gamma_1, \Gamma_2 \subseteq \Gamma$
 $\Delta', \Gamma_2 \mid \Gamma'_1 \vdash E'[e']$ (4), inductive hypothesis with $\Delta' \cap \Gamma_2 = \emptyset$
 $g \notin \Gamma'_1, \text{rng}(S')$ (5), inductive hypothesis
 $\Delta, \Delta' \mid \Gamma - \Gamma_1, \Gamma'_1 \vdash E[e']$ (6), by (4)
 $g \notin \Gamma - \Gamma_1, \Gamma'_1, \text{rng}(S')$ (7), by (2),(5)

We will also need that the tail-context is preserved by evaluation. Here we use the $\mathcal{T}[\bar{f}]$ tail-context, which includes the condition, that $f_i \notin \text{fv}(e_0)$. This is important, since it means that a tail-call can not pass along a function pointer of the same recursive group. If we didn't have that restriction, functions in the same recursive group could call themselves anonymously in non-tail position and the stack size would be unbounded.

Lemma 37. (*Substitution preserves tail contexts*)

If $e = \mathcal{T}[\bar{f}]$ and $\bar{f} \cap \bar{y} = \emptyset$, then $e[\bar{x}:=\bar{y}] = \mathcal{T}[\bar{f}]$.

Proof. By straight-forward induction on $\mathcal{T}[\bar{f}]$, noticing that $\text{fv}(e[\bar{x}:=\bar{y}]) \subseteq \text{fv}(e), \bar{y}$.

Lemma 38. (The eval context is complementary to the tail context)

If $e = \mathcal{T}[\bar{f}]$, $f \in \bar{f}$ and $e = E[f(\bar{y}; \bar{x})]$, then $E = \square$.

If $e = \mathcal{T}[\bar{f}]$, $e = E[e']$ and $E \neq \square$, then $\bar{f} \cap \text{fv}(e') = \emptyset$.

Proof. Same as proof of lemma 30.

Lemma 39. (Store semantics preserves tail contexts)

If $S \mid e \rightarrow_s S' \mid e'$ not using the (*anon_s*) rule, $e = \mathcal{T}[\bar{f}]$ and $\bar{f} \cap \text{rng}(S) = \emptyset$, then $e' = \mathcal{T}[\bar{f}]$.

Proof. We use case analysis on $\mathcal{T}[\bar{f}]$. Let us first assume that $e = e_0$ where $f_i \notin \text{fv}(e_0)$. By induction on the evaluation context E of the evaluation.

Case (*let_s*):

$$\begin{array}{ll} S \mid \text{let } \bar{x} = \bar{y} \text{ in } e \rightarrow_s S \mid e[\bar{x}:=\bar{y}] & \text{definition} \\ \bar{f} \cap \bar{y} = \emptyset, \bar{f} \cap \text{fv}(e) = \emptyset & \text{assumption} \\ \bar{f} \cap \text{fv}(e[\bar{x}:=\bar{y}]) = \emptyset & \text{since } \text{fv}(e[\bar{x}:=\bar{y}]) \subseteq \text{fv}(e), \bar{y} \end{array}$$

Case (*call_s*):

$$\begin{array}{ll} S \mid f(\bar{y}'; \bar{x}') \rightarrow_s S \mid e[\bar{y}:=\bar{y}'; \bar{x}:=\bar{x}'] & \text{for } f(\bar{y}; \bar{x}) = e \in \Sigma \\ \bar{f} \cap \bar{y}, \bar{x} = \emptyset & \text{by definition of the tail context} \\ e = \mathcal{T}[\bar{f}] & \text{by definition of } \Sigma \\ e[\bar{y}:=\bar{y}'; \bar{x}:=\bar{x}'] = \mathcal{T}[\bar{f}] & \text{lemma 37} \end{array}$$

Case (*reuse_s*):

$$\begin{array}{ll} S, \diamond_k \mid C^k x_1 \dots x_k \rightarrow_s S, x \mapsto C^k x_1 \dots x_k \mid x & \text{definition, (fresh } x, k \geq 1) \\ x \notin \bar{f} & \text{since } x \text{ fresh} \end{array}$$

Case (*atom_s*):

$$\begin{array}{ll} S \mid C \rightarrow_s S, x \mapsto C \mid x & \text{definition, (fresh } x) \\ x \notin \bar{f} & \text{since } x \text{ fresh} \end{array}$$

Case (*bmatch_s*):

$$\begin{array}{ll} S, y \mapsto C^k \bar{y} \mid \text{match } y \{ \overline{p \rightarrow e} \} \rightarrow_s S, y \mapsto C^k \bar{y} \mid e_i[\bar{x}:=\bar{y}] & \text{definition, } (p_i = C^k \bar{y}) \\ \bar{y} \cap \bar{f} = \emptyset & \text{since } S \text{ contains no functions} \\ \bar{f} \cap \text{fv}(e_i) = \emptyset & \text{by assumption} \\ \bar{f} \cap \text{fv}(e_i[\bar{x}:=\bar{y}]) = \emptyset & \text{since } \text{fv}(e[\bar{x}:=\bar{y}]) \subseteq \text{fv}(e_i), \bar{y} \end{array}$$

Case (*dmatch_s*):

$$\begin{array}{ll} S, x \mapsto C^k \bar{y} \mid \text{match! } x \{ \overline{p \rightarrow e} \} \rightarrow_s S, \diamond_k \mid e_i[\bar{x}:=\bar{y}] & \text{definition, } (p_i = C^k \bar{x}) \\ \bar{y} \cap \bar{f} = \emptyset & \text{since } S \text{ contains no functions} \\ \bar{f} \cap \text{fv}(e_i) = \emptyset & \text{by assumption} \\ \bar{f} \cap \text{fv}(e_i[\bar{x}:=\bar{y}]) = \emptyset & \text{since } \text{fv}(e[\bar{x}:=\bar{y}]) \subseteq \text{fv}(e_i), \bar{y} \end{array}$$

Case (*dcon_s*):

$$\begin{array}{ll} S, x \mapsto C^k \bar{x} \mid \text{drop } x; e \rightarrow_s S \mid \text{drop } \bar{x}; e & \text{definition} \\ \bar{x} \cap \bar{f} = \emptyset & \text{since } S \text{ contains no functions} \\ \bar{f} \cap \text{fv}(e) = \emptyset & \text{by assumption} \\ \bar{f} \cap \text{fv}(\text{drop } \bar{x}; e) = \emptyset & \text{as above} \end{array}$$

Case ($free_s$):

$$\begin{array}{ll} S, \diamond_k \mid free\ k; e \longrightarrow_s S \mid e & \text{definition} \\ \bar{f} \cap fv(e) = \emptyset & \text{by assumption} \end{array}$$

Case ($eval$):

$$\begin{array}{ll} S \mid E[e] \longrightarrow_s S' \mid E[e'] & (1), \text{ given} \\ \bar{f} \cap fv(E[e]) = \emptyset & (2), \text{ by assumption} \\ S \mid e \longrightarrow_s S' \mid e' & (3), (eval) \text{ rule} \\ \bar{f} \cap fv(e') = \emptyset & (4), \text{ inductive hypothesis} \\ \bar{f} \cap fv(E[e']) = \emptyset & (5), \text{ by (2) and (4)} \end{array}$$

We can now assume that $\mathcal{T}[\bar{f}] \neq e_0$. Then split on outermost layer of the evaluation context E . First we consider the case where $E \neq \square$.

$$\begin{array}{ll} S \mid E[e] \longrightarrow_s S' \mid E[e'] & (1), \text{ given} \\ \bar{f} \cap fv(e) = \emptyset & (2), \text{ by lemma 38} \\ S \mid e \longrightarrow_s S' \mid e' & (3), \text{ by (eval)} \\ \bar{f} \cap fv(e') = \emptyset & (4), \text{ by the first case of this proof} \end{array}$$

We need to show that $E[e']$ is again a tail-context. We split on the remaining cases of E that match a tail-context $\mathcal{T}[\bar{f}] \neq e_0$:

Case $E = g(\bar{y}; \square)$:

$$\begin{array}{ll} E[e] = \mathcal{T}[\bar{f}] & (5), \text{ given} \\ \bar{f} \cap \bar{y} = \emptyset & (6), \text{ by (5)} \\ E[e'] = \mathcal{T}[\bar{f}] & (7), \text{ by (4) and (6)} \end{array}$$

Case $E = \text{let } \bar{x} = \square \text{ in } e_2$:

$$\begin{array}{ll} E[e] = \mathcal{T}[\bar{f}] & (5), \text{ given} \\ e_2 = \mathcal{T}'[\bar{f}] & (6), \text{ by (5)} \\ E[e'] = \mathcal{T}[\bar{f}] & (7), \text{ by (4) and (6)} \end{array}$$

Case $E = \text{match! } \square \{ \bar{p} \mapsto \bar{e} \}$:

$$\begin{array}{ll} E[e] = \mathcal{T}[\bar{f}] & (5), \text{ given} \\ e_i = \mathcal{T}_i[\bar{f}] & (6), \text{ by (5)} \\ E[e'] = \mathcal{T}[\bar{f}] & (7), \text{ by (4) and (6)} \end{array}$$

Lastly, assume that $\mathcal{T}[\bar{f}] \neq e_0$ and $E = \square$. Then split on the cases of $\mathcal{T}[\bar{f}]$.

Case $f_i(\bar{y}; \bar{x})$: Impossible, since only the call rule (which we omitted) can reduce in this case.

Case $\text{let } \bar{x} = \bar{y} \text{ in } \mathcal{T}[\bar{f}]$:

$$\begin{array}{ll} S \mid \text{let } \bar{x} = \bar{y} \text{ in } e \longrightarrow_s S \mid e[\bar{x}:=\bar{y}] & \text{definition} \\ \bar{y} \cap \bar{f} = \emptyset & \text{assumption} \\ e = \mathcal{T}[\bar{f}] & \text{assumption} \\ e[\bar{x}:=\bar{y}] = \mathcal{T}[\bar{f}] & \text{by lemma 37} \end{array}$$

Case $\text{match } y \{ \overline{p_i \mapsto \mathcal{T}_i[\bar{f}]} \}$:

$$\begin{array}{ll} S, y \mapsto C^k \bar{y} \mid \text{match } y \{ \overline{p \mapsto e} \} \longrightarrow_s S, y \mapsto C^k \bar{y} \mid e_i[\bar{x}:=\bar{y}] & \text{definition, } (p_i = C^k \bar{y}) \\ \bar{y} \cap \bar{f} = \emptyset & \text{since } S \text{ contains no functions} \\ e_i = \mathcal{T}[\bar{f}] & \text{assumption} \\ e_i[\bar{x}:=\bar{y}] = \mathcal{T}[\bar{f}] & \text{by lemma 37} \end{array}$$

Case match! $x \{ \overline{p_i \mapsto \mathcal{T}_i[\bar{f}]} \}$:

$$\begin{array}{ll} S, x \mapsto C^k \bar{y} \mid \text{match! } x \{ \overline{p \rightarrow e} \} \longrightarrow_s S, \diamond_k \mid e_i[\bar{x} := \bar{y}] & \text{definition, } (p_i = C^k \bar{x}) \\ \bar{y} \cap \bar{f} = \emptyset & \text{since } S \text{ contains no functions} \\ e_i = \mathcal{T}[\bar{f}] & \text{assumption} \\ e_i[\bar{x} := \bar{y}] = \mathcal{T}[\bar{f}] & \text{by lemma 37} \end{array}$$

Case drop x; $\mathcal{T}[\bar{f}]$:

$$\begin{array}{ll} S, x \mapsto C^k \bar{x} \mid \text{drop } x; e \longrightarrow_s S \mid \text{drop } \bar{x}; e & \text{definition} \\ \bar{x} \cap \bar{f} = \emptyset & \text{since } S \text{ contains no functions} \\ e = \mathcal{T}[\bar{f}] & \text{assumption} \end{array}$$

Case free k; $\mathcal{T}[\bar{f}]$:

$$\begin{array}{ll} S, \diamond_k \mid \text{free } k; e \longrightarrow_s S \mid e & \text{definition} \\ e = \mathcal{T}[\bar{f}] & \text{assumption} \end{array}$$

We can now prove the main lemma of this section:

Lemma 40. (*Second order stack bound*)

Let Σ be non-empty and fully-in-place such that for all functions f in Σ mutually recursive with \bar{f} we have $f(\bar{y}; \bar{x}) = \mathcal{T}[\bar{f}]$. If $e = \mathcal{T}[\bar{f}]$, $\Delta \mid \Gamma \vdash_\Sigma e$ and $|e| \leq |e_{\max}| \cdot |\Sigma|$, and $\Gamma, \text{rng}(S) \cap \text{dom}(\Sigma) = \emptyset$, then at any intermediate step $S \mid e \longrightarrow_s^* S' \mid E'[e']$ we have $|E'| \leq |e_{\max}| \cdot |\Sigma|^2$.

Proof. By induction on Σ . In both the base and the inductive case, we use another induction on all subderivations of $S \mid e \longrightarrow_s^* S' \mid E'[e']$. Our induction hypothesis is then:

- Either $\Sigma = \bar{f}$ or the lemma holds for all derivations on $\Sigma' \subseteq \Sigma$.
- The lemma holds for Σ and all derivations $S_2 \mid e_2 \longrightarrow_s^* S' \mid E'[e']$ such that $S \mid e \longrightarrow_s^* S_2 \mid e_2$, $e_2 = \mathcal{T}[\bar{f}]$, $\Delta' \mid \Gamma' \vdash_\Sigma e_2$, $\Gamma', \text{rng}(S_2) \cap \text{dom}(\Sigma) = \emptyset$ and $|e_2| \leq |e_{\max}| \cdot |\Sigma|$.

Case Base case $S \mid e \longrightarrow_s^* S \mid E'[e']$ with $e = E'[e']$: Clear, since $|E'| \leq |e| \leq |e_{\max}| \cdot |\Sigma|$ and $|\Sigma| \geq 1$.

If $S \mid e \longrightarrow_s^* S' \mid E'[e']$ even without the anon rule, we have

$$\begin{array}{ll} e = \tilde{\mathcal{T}}[\bar{f}] & \text{clear, since } e = \mathcal{T}[\bar{f}] \\ \text{\$D} & \text{G} \\ |e| \leq |e_{\max}| \cdot |\Sigma| & \text{by assumption} \\ |E'| \leq |e_{\max}| \cdot |\Sigma| & \text{by lemma 34} \end{array}$$

Else, we choose the longest subderivation from $E[e]$ that does not involve the anon rule $S \mid e \longrightarrow_s^* S_g \mid E_g[g(\bar{x}_2)]$.

Then:

$$\begin{array}{ll} E_g[g(\bar{x}_2)] = \mathcal{T}[\bar{f}] & \text{by lemma 39} \\ g \notin \bar{f} & \text{by lemma 38} \\ g \in \Sigma - \bar{f} & \text{by above} \end{array}$$

Case If evaluation never leaves g and $E' = E_g[E'']$, then:

FP²: Fully in-Place Functional Programming

$ E_g \leq e_{\max} \cdot \Sigma $	since the derivation involves no (<i>anon_s</i>)
$g(xs_2) = \mathcal{T}[gs]$	definition of $\mathcal{T}[gs]$, where gs is g 's recursive group
$g \mid xs_2 \vdash g(xs_2)$	BAPP
$xs_2, \text{rng}(S_g) \cap \text{dom}(\Sigma) = \emptyset$	repeatedly apply lemma 36
$ g(xs_2) = 1 \leq e_{\max} $	definition of length
$S_g \mid E_g[g(\bar{x}_2)] \longrightarrow_s^* S' \mid E_g[E''[e']]$	since we chose a subderivation
$S_g \mid g(\bar{x}_2) \longrightarrow_s^* S' \mid E''[e']$	by the (<i>eval</i>) rule
$ E'' \leq e_{\max} \cdot \Sigma - \bar{f} \cdot \Sigma - \bar{f} $	by the inductive hypothesis (1)
$ \Sigma - \bar{f} + 1 \leq \Sigma $	definition of $\Sigma - \bar{f}$
$ E' = E_g[E''] = E_g + E'' \leq e_{\max} \cdot \Sigma \cdot \Sigma $	as above

Case If evaluation leaves g and $S_g \mid E_g[g(\bar{x}_2)] \longrightarrow_s S'_g \mid E_g[\bar{x}_3]$, then:

$ E_g[g(xs_2)] \leq e_{\max} \cdot \Sigma $	since the derivation involves no (<i>anon_s</i>)
$ E_g[xs_3] \leq e_{\max} \cdot \Sigma $	substitution of variables does not increase length
$E_g[g(xs_2)] = \mathcal{T}[\bar{f}]$	by lemma 39
$E_g[xs_3] = \mathcal{T}[\bar{f}]$	substitution of variables does not introduce calls
$\Delta' \mid \Gamma' \vdash E_g[xs_3]$	repeatedly apply lemma 33
$\Gamma', \text{rng}(S'_g) \cap \text{dom}(\Sigma) = \emptyset$	repeatedly apply lemma 36
$ E' \leq e_{\max} \cdot \Sigma ^2$	by the inductive hypothesis (2)

Finally, the theorem:

Theorem 8. (*An FIP program uses constant stack space*)

Let Σ be fully in-place such that for all functions f in Σ mutually recursive with \bar{f} we have $f(\bar{y}; \bar{x}) = \mathcal{T}[\bar{f}]$ and let S be a store that contains no functions. Then at any intermediate step $S \mid f(\bar{y}; \bar{x}) \longrightarrow_s^* S' \mid E[e']$, we have $|E| \leq |e_{\max}| \cdot |\Sigma|^2$.

Proof. Apply lemma 40 to $f(\bar{y}; \bar{x})$.

Σ nonempty	since $f(\bar{y}; \bar{x}) = e \in \Sigma$
$f(\bar{y}; \bar{x}) = \mathcal{T}[\bar{f}]$	by assumption
$\bar{y} \mid \bar{x} \vdash f(\bar{y}; \bar{x})$	CALL
$ f(\bar{y}; \bar{x}) = 1 \leq e_{\max} $	definition of length
$\bar{x}, \text{rng}(S) \cap \text{dom}(\Sigma) = \emptyset$	since S contains no functions

$$\begin{aligned}
 f(\bar{y}; \bar{x}) &= e & app(\bar{y}; k, \bar{x}') &= \text{match! } k \\
 f'(\bar{y}; \bar{x}, k) &= \llbracket e \rrbracket_{f,k} & & \quad \text{H} \rightarrow \bar{x}' \\
 & & & \quad Acc_E(\Gamma; k') \rightarrow \llbracket E_G[\bar{x}'] \rrbracket_{f,k'}
 \end{aligned}$$

$$\begin{aligned}
 (tctx) \quad \llbracket \mathbb{T}[e] \rrbracket_{f,k} &= \mathbb{T}[\llbracket e \rrbracket_{f,k}] \\
 (tail) \quad \llbracket E[f(\bar{y}; e_0)] \rrbracket_{f,k} &= f'(\bar{y}; e_0, Acc_E(\Gamma_2; k)) \text{ iff } (\star) \\
 (base) \quad \llbracket e_0 \rrbracket_{f,k} &= app(\bar{y}; k, e_0) \quad \text{where } f \notin \text{fv}(e_0)
 \end{aligned}$$

$$(\star): \text{ for } \Delta, \bar{y} \mid \Gamma_1, \Gamma_2 \vdash E[f(\bar{y}; e)], \text{ we have } \Delta, \bar{y} \mid \Gamma_1 \vdash f(\bar{y}; e) \text{ and } \bar{y} \mid \Gamma_2 \vdash E[\square]$$

$$\begin{aligned}
 \mathbb{T} \quad & ::= \square \mid \text{let } \bar{x} = e_0 \text{ in } \mathbb{T} \mid \text{match } e_0 \{ \overline{p_i \mapsto \mathbb{T}_i} \} \mid \text{match! } e_0 \{ \overline{p_i \mapsto \mathbb{T}_i} \} \\
 & \mid \text{drop } \bar{x}; \mathbb{T} \mid \text{free } k; \mathbb{T} \quad (\text{where } f \notin \text{fv}(e_0))
 \end{aligned}$$

$$\begin{aligned}
 Acc_E(\emptyset; k') &= k' \\
 Acc_E(\Gamma, \diamond_k, x_1, \dots, x_n; k') &= C_E^k x_1 \dots x_n Acc_E(\Gamma; k') C_{unit} \dots C_{unit} \text{ where } 0 \leq n < k
 \end{aligned}$$

Fig. 12. Generalized TRMC with reused contexts

D TRMC WITH REUSE TRANSLATION

In this section, we prove the soundness of the TRMC translation with reused context. We base our formalization on the slightly generalised presentation in figure 12. In particular, the constructors of our zipper are generated by an accumulator function Acc_E , which the version in the paper would always call on \diamond_k, \bar{z}_i with $k = |\bar{z}_i|$. But what if $k \geq |\bar{z}_i| + 1$? We can still represent this by padding the accumulator with atoms C_{unit} . Similarly, if several reuse credits are available, we create a list of accumulators that jointly stores all free variables.

D.1 Soundness

Lemma 41. (*Accumulators are FIP*)

If $Acc_E(\Gamma; k)$ exists, then $\emptyset \mid \Gamma, k \vdash Acc_E(\Gamma; k)$.

Proof. By induction on $Acc_E(\Gamma; k)$:

Case $Acc_E(\emptyset; k) = k$: By the VAR rule we have $\emptyset \mid k \vdash k$.

Case $Acc_E(\Gamma, \diamond_k, x_1, \dots, x_n; k) = C_E^k x_1 \dots x_n Acc_E(\Gamma; k) C_{unit} \dots C_{unit}$:

$\emptyset \mid \Gamma, k \vdash Acc_E(\Gamma; k)$ (1), by inductive hypothesis

$\emptyset \mid \emptyset \vdash C_{unit}$ (2), ATOM

$\emptyset \mid x_i \vdash x_i$ (3), VAR

$\emptyset \mid \Gamma, \diamond_k, x_1, \dots, x_n, k \vdash C_E^k x_1 \dots x_n Acc_E(\Gamma; k) C_{unit} \dots C_{unit}$ (4), REUSE

Lemma 42. (*Matching accumulators is FIP*)

If $Acc_E(\Gamma; k') \neq k'$ exists and $\Delta \mid \Gamma_1, \Gamma_2, k' \vdash e$ then $\Delta \mid \Gamma_1, k \vdash \text{match! } k \{ Acc_E(\Gamma_2; k') \mapsto e \}$.

Proof. We recursively expand $\text{match! } k \{ Acc_E(\Gamma; k') \mapsto e \}$ as:

• $e[k' := k]$ if $Acc_E(\emptyset; k') = k'$

• $\text{match! } k \{ C_E^k x_1 \dots x_n k' u_1 \dots u_m \mapsto \text{drop } u_1; \dots \text{drop } u_m; \text{match! } k' \{ Acc_E(\Gamma; k'') \mapsto e \} \}$ if $Acc_E(\Gamma, \diamond_k, \dots$

Notice that we could also match on the u_i as $\text{match! } u_i C_{unit} \mapsto \dots$ instead of dropping them, to obtain a solution without deallocation.

Then, by induction on $\text{Acc}_E(\Gamma_2; k')$:

Case $\text{Acc}_E(\emptyset; k') = k'$: By assumption we have $\Delta \mid \Gamma_1, \emptyset, k' \vdash e$. Use lemma 11 to obtain $\Delta \mid \Gamma_1, \emptyset, k \vdash e[k' := k]$.

Case $\text{Acc}_E(\Gamma'_2, \diamond_k, x_1, \dots, x_n; k') = C_E^k x_1 \dots x_n \text{Acc}_E(\Gamma'_2; k') C_{\text{unit}} \dots C_{\text{unit}}$:

$\Delta \mid \Gamma'_1, k' \vdash \text{match! } k' \{ \text{Acc}_E(\Gamma'_2; k'') \mapsto e \}$

(1), by inductive hypothesis

$\Delta \mid \Gamma'_1, k', u_1, \dots, u_m \vdash \text{drop } u_1; \dots \text{drop } u_m;$
 $\text{match! } k' \{ \text{Acc}_E(\Gamma'_2; k'') \mapsto e \}$

(2), DROP rule

$\Gamma'_1 = \Gamma_1, \diamond_k, x_1, \dots, x_n$

(3), as $x_i \notin \Gamma'_2$ and $\Delta \mid \Gamma'_1, \Gamma'_2, k' \vdash e$

$\Delta \mid \Gamma_1, k \vdash \text{match! } k \{ C_E^k x_1 \dots x_n k' u_1 \dots u_m \mapsto$

$\text{drop } u_1; \dots \text{drop } u_m; \text{match! } k' \{ \text{Acc}_E(\Gamma; k'') \mapsto e \}$

(4), by MATCH!

$\Gamma_2 := \Gamma'_2, \diamond_k, x_1, \dots, x_n$

(5), define

$\Delta \mid \Gamma_1, k \vdash \text{match! } k \{ \text{Acc}_E(\Gamma_2; k') \mapsto e \}$

(6), by definition

Lemma 43. (*Replacement lemma for tail contexts*)

Let $\Delta \mid \Gamma \vdash \mathbb{T}[e]$ with $\Delta' \mid \Gamma' \vdash e$. For all expressions e' with $\Delta' \mid \Gamma', \Gamma'' \vdash e'$, we have $\Delta \mid \Gamma, \Gamma'' \vdash \mathbb{T}[e']$.

Proof. By induction on E .

Case \square : By assumption, as $\Delta = \Delta', \Gamma = \Gamma'$.

Case let $\bar{x} = e_0$ in \mathbb{T}

$\Delta \mid \Gamma_1, \Gamma_2, \Gamma_3 \vdash \mathbb{T}[e]$

(1), given

$\Delta \mid \Gamma_2, \Gamma_3 \vdash \mathbb{T}'[e]$

(2), by (1) and LET

$\Delta \mid \Gamma_2, \Gamma_3, \Gamma'' \vdash \mathbb{T}'[e']$

(3), by inductive hypothesis

$\Delta \mid \Gamma_1, \Gamma_2, \Gamma_3, \Gamma'' \vdash \mathbb{T}[e']$

(4), by LET

Case match $e_0 \{ \overline{p \mapsto \mathbb{T}} \}$:

$\Delta \mid \Gamma_1, \Gamma_2 \vdash \mathbb{T}[e]$

(1), given

$\Delta, \bar{x}_i \mid \Gamma_2 \vdash \mathbb{T}'[e]$

(2), by (1) and MATCH!

$\Delta, \bar{x}_i \mid \Gamma_2, \Gamma'' \vdash \mathbb{T}'[e']$

(3), by inductive hypothesis

$\Delta \mid \Gamma_1, \Gamma_2, \Gamma'' \vdash \mathbb{T}[e']$

(4), by MATCH!

Case match! $e_0 \{ \overline{p \mapsto \mathbb{T}} \}$:

$\Delta \mid \Gamma_1, \Gamma_2 \vdash \mathbb{T}[e]$

(1), given

$\Delta \mid \Gamma_2, \bar{x}_i \vdash \mathbb{T}'[e]$

(2), by (1) and MATCH!

$\Delta \mid \Gamma_2, \bar{x}_i, \Gamma'' \vdash \mathbb{T}'[e']$

(3), by inductive hypothesis

$\Delta \mid \Gamma_1, \Gamma_2, \Gamma'' \vdash \mathbb{T}[e']$

(4), by MATCH!

Case drop x ; $\mathbb{T} \mid \text{free } k$; \mathbb{T}

$\Delta \mid \Gamma, x \vdash \mathbb{T}[e]$

(1), given

$\Delta \mid \Gamma \vdash \mathbb{T}'[e]$

(2), by (1) and the relevant FIP rule

$\Delta \mid \Gamma, \Gamma'' \vdash \mathbb{T}'[e']$

(3), by inductive hypothesis

$\Delta \mid \Gamma, x, \Gamma'' \vdash \mathbb{T}[e']$

(4), by the relevant FIP rule

Lemma 44. (*The TRMC translation preserves FIP*)

Fix a function $f(\bar{y}; \bar{x})$. If $\Delta, \bar{y} \mid \Gamma \vdash e$, $\llbracket e \rrbracket_{f,k}$ exists and assuming that $\text{app}(\bar{y}; k, \bar{x}')$ is in the signature, then $\Delta, \bar{y} \mid \Gamma, k \vdash \llbracket e \rrbracket_{f,k}$.

Proof. Show that the claim holds for any Δ, Γ, e by induction on the translation rules,

Case (*tctx*).

$$\begin{array}{ll}
 \Delta, \bar{y} \mid \Gamma \vdash \mathbb{T}[e] & \text{given} \\
 \Delta, \bar{y}, \Delta' \mid \Gamma' \vdash e & \text{unwrap } \mathbb{T} \\
 \Delta, \bar{y}, \Delta' \mid \Gamma', k \vdash \llbracket e \rrbracket_{f,k} & \text{inductive hypothesis} \\
 \Delta, \bar{y} \mid \Gamma, k \vdash \mathbb{T}[\llbracket e \rrbracket_{f,k}] & \text{wrap } \mathbb{T} \text{ by lemma 43}
 \end{array}$$

Case (tail).

$$\begin{array}{ll}
 \Delta, \bar{y} \mid \Gamma_1, \Gamma_2 \vdash E[f(\bar{y}; e)] & (1), \text{ given} \\
 \Delta, \bar{y} \mid \Gamma_1 \vdash f(\bar{y}; e) & (2), \text{ given by } (\star) \\
 \bar{y} \mid \Gamma_2 \vdash E[\square] & (3), \text{ given by } (\star) \\
 \Delta, \bar{y} \mid \Gamma_1, k \vdash f'(\bar{y}; e, k) & (4), \text{ by } (2) \\
 \Delta, \bar{y} \mid \Gamma_1, \Gamma_2, k \vdash f'(\bar{y}; e, \text{Acc}_E(\Gamma_2; k)) & (5), \text{ by } (4) \text{ and lemma 41}
 \end{array}$$

Case (base).

$$\begin{array}{ll}
 \Delta, \bar{y} \mid \Gamma \vdash e_0 & (1), \text{ given} \\
 \Delta, \bar{y} \mid \Gamma, k \vdash \text{app}(\bar{y}; k, e_0) & (2), \text{ by assumption}
 \end{array}$$

Lemma 45. (Value replacement lemma)

Let $\Delta \mid \Gamma \vdash E[\square]$ and $\emptyset \mid \Gamma' \vdash \bar{v}$, then $\Delta \mid \Gamma, \Gamma' \vdash E[\bar{v}]$.

Proof. By induction on E .

Case \square : By assumption, as $\Delta, \Gamma = \emptyset$.

Case $C^k x_1 \dots E' \dots v_k \mid (x_1, \dots, E', \dots, v_n) \mid E' e \mid x E' \mid f(\bar{y}; E') \mid \text{match! } E' \{ \overline{p \mapsto e} \}$:

$$\begin{array}{ll}
 \Delta \mid \Gamma \vdash E[\square] & (1), \text{ given} \\
 \Delta \mid \Gamma_i \vdash E'[\square] & (2), \text{ by } (1) \text{ and the relevant FIP rule} \\
 \Delta \mid \Gamma_i, \Gamma' \vdash E'[\bar{v}] & (3), \text{ by inductive hypothesis} \\
 \Delta \mid \Gamma, \Gamma' \vdash E[\bar{v}] & (4), \text{ by the relevant FIP rule}
 \end{array}$$

Case let $\bar{x} = E'$ in e : (special as we add to the borrowed environment)

$$\begin{array}{ll}
 \Delta \mid \Gamma \vdash E[\square] & (1), \text{ given} \\
 \Delta, \Gamma_2 \mid \Gamma_1 \vdash E'[\square] & (2), \text{ by } (1) \text{ and LET} \\
 \Delta, \Gamma_2 \mid \Gamma_1, \Gamma' \vdash E'[\bar{v}] & (3), \text{ by inductive hypothesis} \\
 \Delta \mid \Gamma, \Gamma' \vdash E[\bar{v}] & (4), \text{ by LET}
 \end{array}$$

Lemma 46. (The app function is FIP.)

Fix a function $f(\bar{y}; \bar{x})$. If $\Delta, \bar{y} \mid \Gamma \vdash e$ and $\llbracket e \rrbracket_{f,k}$ exists, then $\bar{y} \mid \bar{x}', k \vdash \text{app}(\bar{y}; k, \bar{x}')$.

Proof. We have:

$$\begin{array}{l}
 \text{app}(\bar{y}; k, \bar{x}') = \text{match! } k \\
 \quad \text{H} \rightarrow \bar{x}' \\
 \quad \text{Acc}_E(\Gamma; k') \rightarrow \llbracket E_G[\bar{x}'] \rrbracket_{f,k'}
 \end{array}$$

Apply the **MATCH!** rule.

Case H: Clearly $\emptyset \mid \bar{x}' \vdash \bar{x}'$.

Case $\text{Acc}_E(\Gamma)$:

$$\begin{array}{ll}
 \bar{y} \mid \Gamma \vdash E_G[\square] & (1), \text{ by } (\star) \text{ rule} \\
 \bar{y} \mid \Gamma, \bar{x}' \vdash E_G[\bar{x}'] & (2), \text{ by lemma 45} \\
 \bar{y} \mid \Gamma, \bar{x}', k' \vdash \llbracket E_G[\bar{x}'] \rrbracket_{f,k'} & (3), \text{ by lemma 44} \\
 \bar{y} \mid \bar{x}', k \vdash \text{match! } k \{ \text{Acc}_E(\Gamma; k') \mapsto \llbracket E_G[\bar{x}'] \rrbracket_{f,k'} \} & (4), \text{ by lemma 42}
 \end{array}$$

Theorem 9. (*The TRMC transformation is sound.*)

Let f be a function with $\bar{y} \mid \bar{x} \vdash f(\bar{y}; \bar{x})$ and $f(\bar{v}_1; \bar{v}_2) \longrightarrow^* \bar{w}$. If it can be transformed into f' , then $\bar{y} \mid \bar{x}, k \vdash f'(\bar{y}; \bar{x}, k)$, $\bar{y} \mid k, \bar{x} \vdash \text{app}(\bar{y}; k, \bar{x})$ and $f'(\bar{v}_1; \bar{v}_2, H) \longrightarrow^* \bar{w}$.

Proof. Our transformation follows the general framework of TRMC where we use a defunctionalized evaluation context (section 4.2 of [Leijen and Lorenzen 2023]). Note that we generalize the (*tlet*) and (*tmatch*) rules of the TRMC translation (figure 2 in [Leijen and Lorenzen 2023]) to a (*tctx*) rule. In particular, we include drops and frees in the tail context. However, the correctness of these additional cases is clear. We thus have $f'(\bar{v}_1; \bar{v}_2, H) \longrightarrow^* \bar{w}$ by the correctness of the context transformation and the context laws for defunctionalized evaluation contexts.

We have $\bar{y} \mid \bar{x}, k \vdash f'(\bar{y}; \bar{x}, k)$ by lemma 44 and $\bar{y} \mid k, \bar{x}' \vdash \text{app}(\bar{y}; k, \bar{x}')$ by lemma 46.

E SOUNDNESS OF HEAP SEMANTICS

This section contains the soundness result for the heap semantics. Since the heap semantics extends the store semantics, we will follow the same proof structure. In particular, many lemmata hold unchanged (just replacing \longrightarrow_s^* by \longrightarrow_h^* , S with H , using \otimes , etc). Thus we only add the extra cases to each lemma. We extend the free variables as usual:

$$\begin{aligned} \text{fv}(\text{dup } \bar{x}; e) &:= \text{fv}(e), \bar{x} \\ \text{fv}(\text{dropru } \bar{x}; e) &:= \text{fv}(e), \bar{x} \\ \text{fv}(\text{alloc } k; e) &:= \text{fv}(e) \\ \text{fv}(\lambda^{\bar{z}} \bar{x}. e) &:= \text{fv}(e) - \bar{x} \end{aligned}$$

E.1 Properties of Heaps

As before, we will often need to focus on a specific part of the heap that corresponds to all values reachable from a root set Γ . In this section, we generalize the properties of stores to heaps using the join operator:

$$\begin{aligned} \emptyset \otimes H_2 &= H_2 \\ H_1, \diamond_k \otimes H_2 &= H_1 \otimes H_2, \diamond_k \\ H_1, x \mapsto^n v \otimes H_2 &= H_1 \otimes H_2, x \mapsto^n v \quad \text{iff } x \notin \text{dom}(H_2) \\ H_1, x \mapsto^n v \otimes H_2, x \mapsto^m v, z \mapsto^{k+1} w &= H_1 \otimes H_2, x \mapsto^{n+m} v, z \mapsto^k w \quad \text{iff } \bar{z} = \text{fv}(v) \end{aligned}$$

This definition still works if H_1 contains cycles. We can move the elements of the cycle in any order into H_2 . It might be attractive to decrease the counts of \bar{z} in H_1 itself, since this would correspond to “deleting” the element x from H_1 . However, with that modification, we could not handle cycles, because then the children \bar{z} of the last cycle element would be in H_2 instead of H_1 .

We write $\#^x \Gamma$ to denote the number of times x occurs in Γ , $\#^x H_1$ to denote the number of times it occurs free in $\text{rng}(H_1)$ and $\#^x (H_1 \cap H_2)$ to denote the number of times it occurs free in $\text{rng}(H_1) \cap \text{rng}(H_2)$.

Lemma 47. (*Counts in joined heaps*)

We have $\#^x (H_1 \otimes H_2) = \#^x H_1 + \#^x H_2 - \#^x (H_1 \cap H_2)$.

Proof. Since the join operator merges compatible $v \in \text{rng}(H_1) \cap \text{rng}(H_2)$.

Lemma 48. (*Distributivity of counting*)

We have $\#^x ((H_1 \otimes H_2) \cap H_3) = \#^x (H_1 \cap H_3) + \#^x (H_2 \cap H_3) - \#^x (H_1 \cap H_2 \cap H_3)$.

Proof. By lemma 47 on those v also in H_3 .

Lemma 49. (*Joining heaps*)

Let H_1 be a sound/linear heap with roots Γ_1 and H_2 a linear heap with roots Γ_2 . Then $H_1 \otimes H_2$ is a sound/linear heap with roots Γ_1, Γ_2 .

Proof.

Case Well-definedness: The reference counts of the \bar{z} can become zero temporarily. The count can not become negative, because we only decrease the count of \bar{z} if their parent is also in H_2 . This can only happen as often as the \bar{z} occur free in $\text{rng}(H_2)$. Since H_2 is linear, the count of \bar{z} is at least the number of times \bar{z} occurs free in $\text{rng}(H_2)$. Furthermore, since \bar{z} has a reference count of at least one in H_1 , we will end up with a positive reference count in the end.

Case Soundness: The invariant that $x \in \text{dom}(H_1 \otimes H_2)$ iff $x \in \text{dom}(H_1)$ or $x \in \text{dom}(H_2)$ is maintained in all cases of the definition \otimes . Since H_1 and H_2 are sound, then so is $H_1 \otimes H_2$.

Case Linearity: Assume that H_1 is linear and $x \mapsto^n v \in H_1$, $x \mapsto^m v \in H_2$. The final reference count of x is $n + m - \#^x (H_1 \cap H_2)$.

$$\begin{aligned}
 n &= \#^x \Gamma_1 + \#^x H_1 && (1), \text{ since } H_1 \text{ linear} \\
 m &= \#^x \Gamma_2 + \#^x H_2 && (2), \text{ since } H_2 \text{ linear} \\
 n + m - \#^x (H_1 \cap H_2) &&& (3), \text{ final reference count} \\
 &= \#^x \Gamma_1 + \#^x \Gamma_2 + \#^x H_1 + \#^x H_2 - \#^x (H_1 \cap H_2) && (4), \text{ by (1) and (2)} \\
 &= \#^x \Gamma_1 + \#^x \Gamma_2 + \#^x (H_1 \otimes H_2) && (5), \text{ by lemma 47}
 \end{aligned}$$

This shows that Γ_1, Γ_2 are the roots of $H_1 \otimes H_2$.

Lemma 50. (*Associativity of joining heaps*)

Let H_1 be a sound heap and H_2, H_3 be linear heaps. Then $(H_1 \otimes H_2) \otimes H_3 = H_1 \otimes (H_2 \otimes H_3)$.

Proof. As before $x \in \text{dom}((H_1 \otimes H_2) \otimes H_3)$ iff $x \in \text{dom}(H_1)$, $x \in \text{dom}(H_2)$ or $x \in \text{dom}(H_3)$ iff $x \in \text{dom}(H_1 \otimes (H_2 \otimes H_3))$. Let us consider the reference count of one such x with reference count n, m, k in H_1, H_2, H_3 respectively. Then:

$$\begin{aligned}
 n + m + k - \#^x (H_1 \cap H_2) - \#^x ((H_1 \otimes H_2) \cap H_3) &&& (1), \text{ reference count in } (H_1 \otimes H_2) \otimes H_3 \\
 &= n + m + k - \#^x (H_1 \cap H_2) - \#^x (H_1 \cap H_3) \\
 &\quad - \#^x (H_2 \cap H_3) + \#^x (H_1 \cap H_2 \cap H_3) && (2), \text{ lemma 48} \\
 &= n + m + k - \#^x (H_1 \cap (H_2 \otimes H_3)) - \#^x (H_2 \cap H_3) && (3), \text{ reference count in } H_1 \otimes (H_2 \otimes H_3)
 \end{aligned}$$

Lemma 51. (*Symmetry of joining heaps*)

Let H_1, H_2 be linear heaps. Then $H_1 \otimes H_2 = H_2 \otimes H_1$.

Proof. Clearly, $\text{dom}(H_1 \otimes H_2) = \text{dom}(H_1), \text{dom}(H_2) = \text{dom}(H_2), \text{dom}(H_1) = \text{dom}(H_2 \otimes H_1)$. By lemma 47, the reference counts are the same.

We can define a category of linear, compatible heaps, where there is a unique morphism $f : H_1 \rightarrow H_2$ between H_1, H_2 iff $\text{dom}(H_1) \subseteq \text{dom}(H_2)$ and $\text{roots}(H_1) \subseteq \text{roots}(H_2)$. In that case we just write $H_1 \rightarrow H_2$.

Lemma 52. (*Monotocity of joining heaps*)

Let $H_1 \rightarrow H_2$ and $H_3 \rightarrow H_4$. Then $H_1 \otimes H_3 \rightarrow H_2 \otimes H_4$.

Proof. As before, the variables of $H_1 \otimes H_3$ are those of H_1, H_3 . Those variables occur in H_2, H_4 and thus $H_2 \otimes H_4$. By lemma 49, the roots of $H_1 \otimes H_3$ is the disjoint union of the roots of H_1 and H_3 . By assumption, these are a subset of the roots of H_2 and H_4 , which give the roots of $H_2 \otimes H_4$.

This shows that the category is monoidal with the join operator as the tensor product and the empty heap as tensor unit. For a linear heap with $\Gamma \subseteq \text{roots}(H)$, we write $H[\Gamma]$ for the smallest linear subset of H containing Γ . Next we want to define a heap subtraction operator, which will become the internal hom:

$$\begin{aligned}
 [(H_1, \diamond_k), (H_2, \diamond_k)] &= [H_1, H_2] \\
 [(H_1, x \mapsto^n v), (H_2, x \mapsto^n v)] &= [H_1, H_2] \quad \text{iff } x \notin \text{fv}(\text{rng}(H_2)) \\
 [H_1, H_2] &= \llbracket H_1, H_2 \rrbracket \quad \text{iff } \text{dom}(H_1) \subseteq \text{fv}(\text{rng}(H_2))
 \end{aligned}$$

$$\llbracket \emptyset, H_2 \rrbracket = H_2$$

$$\llbracket (H_1, x \mapsto^n v), (H_2, x \mapsto^{n+m} v, z \mapsto^k w) \rrbracket = \llbracket H_1, (H_2, x \mapsto^m v, z \mapsto^{k+1} w) \rrbracket \quad \text{iff } \bar{z} = \text{fv}(v)$$

Lemma 53. (*Heap subtraction*)

If H_1 and H_2 are linear heaps with roots Γ_1, Γ_2 and $H_1 \rightarrow H_2$, then $[H_1, H_2]$ is a linear heap with roots $\Gamma_2 - \Gamma_1$.

Proof.

Case Well-definedness: The operator acts in two stages. In the first stage, all reuse credits are removed from H_1 . Since $\text{roots}(H_1) \subseteq \text{roots}(H_2)$, each reuse credit in H_1 has a matching one in H_2 . Then we recursively remove all $x \in \text{dom}(H_1)$ which do not occur free in $\text{rng}(H_2)$. Since $H_1 \rightarrow H_2$, any $x \in \text{dom}(H_1)$ is also in $\text{dom}(H_2)$ with no bigger reference count. In the second stage, we adjust the reference counts of all other such x .

Case Soundness: H_2 is sound by assumption and we only delete x from H_2 if it is not free in $\text{rng}(H_2)$.

Case Linearity: Clearly, $[H_1, H_2]$ contains all reuse tokens of H_2 not contained in H_1 . Let us focus on the reference counts. After the first stage, H_1 and H_2 are still linear, even if both have new roots. In the second stage, the reference count of any $x \in \text{dom}(H_2)$ becomes:

$$\begin{aligned}
 n &= \#^x \Gamma_1 + \#^x H_1 && (1), \text{ since } H_1 \text{ linear} \\
 n + m &= \#^x \Gamma_2 + \#^x H_2 && (2), \text{ since } H_2 \text{ linear} \\
 m + \#^x (H_1 \cap [H_1, H_2]) &&& (3), \text{ final reference count} \\
 &= \#^x \Gamma_2 + \#^x H_2 - \#^x \Gamma_1 - \#^x H_1 + \#^x (H_1 \cap [H_1, H_2]) && (4), \text{ by (1) and (2)} \\
 &= \#^x \Gamma_2 + \#^x H_2 - \#^x \Gamma_1 + \#^x [H_1, H_2] - \#^x (H_1 \otimes [H_1, H_2]) && (5), \text{ by lemma 47} \\
 &= \#^x \Gamma_2 - \#^x \Gamma_1 + \#^x [H_1, H_2] && (6), \text{ by lemma 57}
 \end{aligned}$$

Lemma 54. (Hom isomorphism)

If $H_2 \rightarrow H_3$, then $H_1 \otimes H_2 \rightarrow H_3$ iff $H_1 \rightarrow [H_2, H_3]$.

Proof.

Case Domain: We have $\text{dom}([H_2, H_3]) \subseteq \text{dom}(H_3)$. If $\text{dom}(H_1) \subseteq \text{dom}([H_2, H_3])$ and $\text{dom}(H_2) \subseteq \text{dom}(H_3)$, then $\text{dom}(H_1), \text{dom}(H_2) \subseteq \text{dom}(H_3)$.

If $H_1 \otimes H_2 \rightarrow H_3$, then not only $\text{dom}(H_1), \text{dom}(H_2) \subseteq \text{dom}(H_3)$, but in particular any $x \in \text{dom}(H_3)$ will only be deleted in the first phase if not in $\text{dom}(H_1)$: Deletion requires that x and all its parents have the same reference count in H_2 and H_3 , and that x or one of its parents is a root in H_3 . But then x or one of its parents will also be a root in H_1 , so the reference count would have to be higher in H_3 than in H_2 . Thus $\text{dom}(H_1) \subseteq \text{dom}([H_2, H_3])$.

Case Roots: Let $\Gamma_1, \Gamma_2, \Gamma_3$ be the roots of H_1, H_2, H_3 respectively. Clearly, $\Gamma_1, \Gamma_2 \subseteq \Gamma_3$ iff $\Gamma_1 \subseteq \Gamma_3 - \Gamma_2$. By lemma 49 the first term gives the roots of $H_1 \otimes H_2$ and by lemma 53 the second term gives the roots of $[H_2, H_3]$.

Lemma 55. (Monotocity of subtracting heaps)

If $H_3 \rightarrow H_1 \rightarrow H_2 \rightarrow H_4$ then $[H_1, H_2] \rightarrow [H_3, H_4]$.

Proof.

Case Domain: The domain of $[H_3, H_4]$ consists of all variables in H_4 which are not deleted. Any such variable will be a root of or reachable from a root in H_3 . Since $H_3 \rightarrow H_1$, the reference counts of H_1 are at least as big as those in H_3 . Thus any deletion in the construction of $[H_3, H_4]$ will also happen during the construction of $[H_1, H_2]$.

Case Roots:

$$\begin{aligned}
 &\text{roots}([H_1, H_2]), \text{roots}(H_1) = \text{roots}(H_2) && (1), \text{ by lemma 53} \\
 &\subseteq \text{roots}(H_4) = \text{roots}([H_3, H_4]), \text{roots}(H_3) && (2), \text{ by lemma 53} \\
 &\text{roots}(H_3) \subseteq \text{roots}(H_1) && (3), \text{ by assumption} \\
 &\text{roots}([H_1, H_2]) \subseteq \text{roots}([H_3, H_4]) && (4), \text{ by (2) and (3)}
 \end{aligned}$$

Notice that this does not quite give us a closed monoidal category, since the internal hom $[H_1, H_2]$ is only defined if $H_1 \rightarrow H_2$. Nonetheless, many results about closed monoidal categories also hold in

our setting. As an example of the usefulness of the categorical perspective, we show that subtracting two heaps one-by-one yields the same result as subtracting them as one joined heap:

Lemma 56. (*Internalized hom isomorphism*)

If $H_2 \rightarrow H_3$ and $H_1 \rightarrow [H_2, H_3]$, then $[H_1 \otimes H_2, H_3] = [H_1, [H_2, H_3]]$.

Proof. Let H be a heap with $H \rightarrow [H_1 \otimes H_2, H_3]$. Then:

- $H \rightarrow [H_1 \otimes H_2, H_3]$ (1), by assumption
- iff $(H \otimes (H_1 \otimes H_2)) \rightarrow H_3$ (2), by lemma 54
- iff $((H \otimes H_1) \otimes H_2) \rightarrow H_3$ (3), by lemma 50
- iff $(H \otimes H_1) \rightarrow [H_2, H_3]$ (4), by lemma 54
- iff $H \rightarrow [H_1, [H_2, H_3]]$ (5), by lemma 54

By instantiating H to $[H_1 \otimes H_2, H_3]$ and $[H_1, [H_2, H_3]]$ using reflexivity, we obtain the equality.

And we can show the existence of the universal morphisms:

Lemma 57. (*Existence of counit*)

If $H_1 \rightarrow H_2$ then $[H_1, H_2] \otimes H_1 = H_2$.

Proof. $[H_1, H_2] \otimes H_1 \rightarrow H_2$: By lemma 54, we have $[H_1, H_2] \otimes H_1 \rightarrow H_2$ iff $[H_1, H_2] \rightarrow [H_1, H_2]$, which is true by reflexivity.

$H_2 \rightarrow [H_1, H_2] \otimes H_1$: We have $\text{dom}(H_2) \subseteq (\text{dom}(H_2) - \text{dom}(H_1)), \text{dom}(H_1) \subseteq \text{dom}([H_1, H_2]), \text{dom}(H_1)$. Additionally, $\text{roots}(H_2) = (\text{roots}(H_2) - \text{roots}(H_1)), \text{roots}(H_1)$.

Lemma 58. (*Existence of unit*)

For any heaps H_1, H_2 we have $H_1 \rightarrow [H_2, H_1 \otimes H_2]$.

Proof. By lemma 54, we have $H_1 \rightarrow [H_2, H_1 \otimes H_2]$ iff $H_1 \otimes H_2 \rightarrow H_1 \otimes H_2$, which is true by reflexivity.

Notice that the above inequality is strict since $[H_2, H_1 \otimes H_2]$ contains all cycles from H_2 . However, the cycles do not matter too much, since they will also be present in all other heaps which H_2 was subtracted from:

Lemma 59. (*Universal morphism for unit*)

If $H_2 \rightarrow H_3$ and $H_1 \rightarrow [H_2, H_3]$ then $[H_2, H_1 \otimes H_2] \rightarrow [H_2, H_3]$.

Proof. We have $H_1 \otimes H_2 \rightarrow H_3$ by lemma 54, which implies the claim by lemma 55.

Our last lemma says that if it is possible to subtract H_1 from H_2 , we can add any H_3 equally before or after the subtraction.

Lemma 60. (*Joining distributes over subtraction*)

If $H_1 \rightarrow H_2$ then $[H_1, H_2 \otimes H_3] = [H_1, H_2] \otimes H_3$.

Proof.

Case Domain: Since $H_1 \rightarrow H_2$, an $x \in \text{dom}(H_3)$ can not get deleted in the construction of $[H_1, H_2 \otimes H_3]$. Thus $\text{dom}([H_1, H_2 \otimes H_3]) = \text{dom}([H_1, H_2]), \text{dom}(H_3)$.

Case Roots:

- $\text{roots}([H_1, H_2 \otimes H_3])$ (1)
- $= (\text{roots}(H_2), \text{roots}(H_3) - \text{roots}(H_1))$ (2)
- $= (\text{roots}(H_2) - \text{roots}(H_1), \text{roots}(H_3))$ (3)
- $= \text{roots}([H_1, H_2] \otimes H_3)$ (4)

E.2 Simple Invariant

Our simple invariant maintains that the heap remains “well-formed” during evaluation. We use two heaps: a “borrowed” heap H_1 which is unchanged by evaluation and an “owned” heap which can be changed. Our simple invariant $I(e, \Delta, \Gamma, H_1, H_2)$ is defined as:

- $\Delta \mid \Gamma \vdash e$
- H_1, H_2 are compatible heaps
- $\Delta \subseteq \text{dom}(H_1)$ and H_1 sound
- $\Gamma = \text{roots}(H_2)$ and H_2 linear

That invariant makes it safe to modify H_2 , as all values that the heap semantics destroys are in Γ and not used anywhere else in the heap ($\Gamma = \text{roots}(H_2)$). It would be enough for soundness to demand $\Gamma \subseteq \text{roots}(H_2)$. However, the stronger assertion directly gives us the garbage-free theorem and by the weakening lemma 64 we can always add separated memory to the heap later.

Unlike in the store semantics, the weaking lemma does not hold for every step. If we have an expression like `drop x`, this will lead to different steps depending on whether the reference count is one or not. Therefore, we need to characterize dropping and dopping in terms of the final heap:

Lemma 61. (*Dupping joins a root*)

If $H \mid \text{dup } x; e \longrightarrow_h H' \mid e$ then $H' = H \otimes H[x]$.

Proof. We have $H, x \mapsto^n v \mid \text{dup } x; e \longrightarrow_h H, x \mapsto^{n+1} v \mid e$. Since $\text{dom}(H \otimes H[x]) = \text{dom}(H)$ and $\text{roots}(H \otimes H[x]) = \text{roots}(H), x$, we have the claim.

Lemma 62. (*Dropping subtracts a root*)

If x is a root of H and $H \mid \text{drop } x; e \longrightarrow_h^* H' \mid e$ then $H' = [H[x], H]$.

Proof. We show the statement for all x by induction over all H with $H[x] \rightarrow H$. In the inductive hypothesis, we can assume the claim for all y and $H' \subsetneq H$ with $H'[y] \rightarrow H'$.

Case If the reference count of x in H is one, then $H_1, x \mapsto^1 v \mid \text{drop } x; e \longrightarrow_h H_1 \mid \text{drop } \bar{z}; e$. Since $H_1 \subsetneq H$ and the \bar{z} are roots of H_1 , we can apply the inductive hypothesis to obtain $H' = [H_1[z_1], \dots, [H_1[z_n], H_1]]$. Since $H[x] = (H_1[z_1] \otimes \dots \otimes H_1[z_n])$, $x \mapsto^1 v$ this implies the claim.

Case If the reference count of x in H is bigger than one, then $H_1, x \mapsto^{n+1} v \mid \text{drop } x; e \longrightarrow_h H_1, x \mapsto^n v \mid e$. Since $[H[x], H] = \llbracket H[x], H \rrbracket$, every variable except x will have an unchanged count and only the count of x will be decreased by one.

Lemma 63. (*Dropru subtracts a root and adds a credit*)

If $H \mid \text{dropru } x; e \longrightarrow_h^* H' \mid e$ then $H' = [H[x], H] \otimes \diamond_k$.

Proof. The rule for `dropru` acts like the rule for `drop`, except that it puts a space credit into the heap. The claim thus follows from lemma 62.

Then we can weaken an evaluation if it completes all dropping procedures: we over-approximate this by requesting that the final e' does not start with a drop. The interesting part of the proof is that we can move H_1 out of the joins and subtractions induced by dups/drops by lemma 50 and 60.

Lemma 64. (*Weakening for heap semantics*)

If $H \mid e \mapsto_h^* H' \mid e'$ and $e' \neq E[\text{drop } x; e'']$, then $H \otimes H_1 \mid e \mapsto_h^* H' \otimes H_1 \mid e'$ for any compatible H_1 .

Proof. By induction on the judgement $H \mid e \mapsto_h^* H' \mid e'$. The reflexive case is obvious. Assume that $H \mid e \mapsto_h^* H' \mid e' \mapsto_h H'' \mid e''$ and $H \otimes H_1 \mid e \mapsto_h^* H' \otimes H_1 \mid e'$. Use case analysis on the last step. Then $H' \otimes H_1 \mid e' \mapsto_h H'' \otimes H_1 \mid e''$ holds trivially for the cases which do not modify

FP²: Fully in-Place Functional Programming

the heap: (let_h), ($call_h$), ($bmatch_h$), (app_h).

Case ($atom_h$, $reuse_h$).

$$\begin{aligned}
 H' \mid E[C^k x_1 \dots x_k] &\mapsto_h [\diamond_k, H'], x \mapsto^1 C^k x_1 \dots x_k \mid E[x] && (1), \text{ by assumption} \\
 [\diamond_k, H' \otimes H_1], x \mapsto^1 C^k x_1 \dots x_k & && \\
 = ([\diamond_k, H'] \otimes H_1), x \mapsto^1 C^k x_1 \dots x_k & && (2), \text{ by lemma 60} \\
 = ([\diamond_k, H'], x \mapsto^1 C^k x_1 \dots x_k) \otimes H_1 & && (3), \text{ since } x \notin \text{dom}(H_1)
 \end{aligned}$$

Case (lam_h).

$$\begin{aligned}
 H' \mid E[\lambda \bar{x}. e] &\mapsto_h [\diamond_k, H'], x \mapsto^1 \lambda \bar{x}. e \mid E[x] && (1), \text{ by assumption} \\
 [H' \otimes H_1], x \mapsto^1 \lambda \bar{x}. e & && \\
 = (H', x \mapsto^1 \lambda \bar{x}. e) \otimes H_1 & && (2), \text{ since } x \notin \text{dom}(H_1)
 \end{aligned}$$

Case ($alloc_h$).

$$\begin{aligned}
 H' \mid E[\text{alloc } k; e] &\mapsto_h H' \otimes \diamond_k \mid E[e] && (1), \text{ by assumption} \\
 H' \otimes H_1 \otimes \diamond_k &= H' \otimes \diamond_k \otimes H_1 && (2), \text{ by lemma 51}
 \end{aligned}$$

Case ($free_h$).

$$\begin{aligned}
 H' \mid E[\text{free } k; e] &\mapsto_h [\diamond_k, H'] \mid E[e] && (1), \text{ by assumption} \\
 [\diamond_k, H' \otimes H_1] &= [\diamond_k, H'] \otimes H_1 && (2), \text{ by lemma 60}
 \end{aligned}$$

Case (dup_h).

$$\begin{aligned}
 H' \mid E[\text{dup } x; e] &\mapsto_h H' \otimes H'[x] \mid E[e] && (1), \text{ by lemma 61} \\
 H' \otimes H_1 \otimes (H' \otimes H_1)[x] & && \\
 = H' \otimes H_1 \otimes H'[x] & && (2), \text{ since } x \in \text{dom}(H') \\
 = H' \otimes H'[x] \otimes H_1 & && (3), \text{ by lemma 51}
 \end{aligned}$$

Case ($dcon_h$), ($dlam_h$), ($drop_h$).

$$\begin{aligned}
 H' \mid E[\text{drop } x; e] &\mapsto_h [H'[x], H'] \mid E[e] && (1), \text{ by lemma 62} \\
 [(H' \otimes H_1)[x], H' \otimes H_1] & && \\
 = [H'[x], H' \otimes H_1] & && (2), \text{ since } x \in \text{dom}(H') \\
 = [H'[x], H'] \otimes H_1 & && (3), \text{ by lemma 60}
 \end{aligned}$$

Case ($dconru_h$), ($dropru_h$).

$$\begin{aligned}
 H' \mid E[\text{dropru } x; e] &\mapsto_h [H'[x], H'] \otimes \diamond_k \mid E[e] && (1), \text{ by lemma 63} \\
 [(H' \otimes H_1)[x], H' \otimes H_1] \otimes \diamond_k & && \\
 = [H'[x], H' \otimes H_1] \otimes \diamond_k & && (2), \text{ since } x \in \text{dom}(H') \\
 = [H'[x], H'] \otimes H_1 \otimes \diamond_k & && (3), \text{ by lemma 60} \\
 = [H'[x], H'] \otimes \diamond_k \otimes H_1 & && (4), \text{ by lemma 51}
 \end{aligned}$$

E.3 Progress

In this section we want to show that the heap semantics can progress if the simple invariant is true and operational semantics can progress. We assume throughout that for any function $f(\bar{y}; \bar{x}) = e \in \Sigma$, we have $\text{fv}(e) \subseteq \bar{y}, \bar{x}$. This is true if Σ is fully in-place:

Lemma 65. (*Free variables of FIP expressions are in Δ, Γ*)

If $\Delta \mid \Gamma \vdash e$, then $\text{fv}(e) \subseteq \Delta, \Gamma$.

Proof. By induction on the judgement $\Delta \mid \Gamma \vdash e$.

Case DUP:

$\Delta \mid \Gamma \vdash \text{dup } x; e$ (1), given
 $\Delta \mid \Gamma, x \vdash e$ (2), by DUP
 $x \in (\Delta, \Gamma)$ (3), by DUP
 $\text{fv}(e) \subseteq \Delta, \Gamma, x$ (4), inductive hypothesis
 $\text{fv}(\text{dup } x; e) \subseteq \Delta, \Gamma$ (5), by (3) and (4)

Case DROPRU:

$\Delta \mid \Gamma, x \vdash \text{dropru } x; e$ (1), given
 $\Delta \mid \Gamma, \diamond_k \vdash e$ (2), by DROPRU
 $\text{fv}(e) \subseteq \Delta, \Gamma, \diamond_k$ (3), inductive hypothesis
 $\text{fv}(\text{dropru } x; e) \subseteq \Delta, \Gamma, x$ (4), by (3)

Case ALLOC:

$\Delta \mid \Gamma \vdash \text{alloc } k; e$ (1), given
 $\Delta \mid \Gamma, \diamond_k \vdash e$ (2), by ALLOC
 $\text{fv}(e) \subseteq \Delta, \Gamma, \diamond_k$ (3), inductive hypothesis
 $\text{fv}(\text{alloc } k; e) \subseteq \Delta, \Gamma$ (4), by (3)

Case MATCH:

$\Delta \mid \Gamma \vdash \text{match } x \{ C_i \bar{x}_i \mapsto \text{dup } \bar{x}_i; e_i \}$ (1), given
 $\Delta \mid \Gamma, \bar{x}_i \vdash e_i$ (2), by MATCH
 $x \in (\Delta, \Gamma)$ (3), by MATCH
 $\text{fv}(e_i) \subseteq \Delta, \Gamma, \bar{x}_i$ (4), inductive hypothesis
 $\text{fv}(\text{match } x \{ C_i \bar{x}_i \mapsto e_i \}) = x, \text{fv}(e_1) - \bar{x}_1, \dots, \text{fv}(e_n) - \bar{x}_n \subseteq \Delta, \Gamma$ (5), definition and (3)

Case APP:

$\Delta \mid \Gamma_1, \Gamma_2 \vdash e_1 e_2$ (1), given
 $\Delta \mid \Gamma_1 \vdash e_1$ (2), by APP
 $\Delta \mid \Gamma_2 \vdash e_2$ (3), by APP
 $\text{fv}(e_1), \text{fv}(e_2) \subseteq \Delta, \Gamma$ (4), inductive hypothesis
 $\text{fv}(e_1 e_2) = \text{fv}(e_1), \text{fv}(e_2) \subseteq \Delta, \Gamma$ (5), definition

Case LAM:

$\Delta \mid \bar{z} \vdash \lambda^{\bar{z}} \bar{x}. e$ (1), given
 $\emptyset \mid \bar{z}, \bar{x} \vdash e$ (2), by LAM
 $\text{fv}(e) \subseteq \bar{x}, \bar{z}$ (3), inductive hypothesis
 $\text{fv}(\lambda^{\bar{z}} \bar{x}. e) \subseteq \bar{z}$ (4), definition

We define $[H - \bar{x}]e$ as the substitution which replaces every variable $y \in \text{fv}(e) - \bar{x}$ by $[H]y$. Notice that this is different from formally subtracting \bar{x} from the heap ($[[H[\bar{x}], H]e$), since the substitution $[H - \bar{x}]e$ can still work on children of \bar{x} , which might get deleted in $[H[\bar{x}], H]$. Then:

Lemma 66. (*Heap Substitution on unused variables*)

If $x \notin \text{fv}(e)$ or $x \notin \text{dom}(H)$, then $[H]e = [H - x]e$.

Proof. As before

Lemma 67. (*Heap Substitution commutes*)

If H sound and $[H]\bar{z} = \bar{v}$, then $[H](e[\bar{x}:=\bar{z}]) = ([H - \bar{x}]e)[\bar{x}:=\bar{v}]$.

Proof. As before

Lemma 68. (*Heap semantics reads values*)

If $I(\bar{v}, \Delta, \Gamma, H_1, H_2)$ then $H_1 \otimes H_2 \mid \bar{v} \xrightarrow{*}_h H_1 \otimes H_2 \mid \bar{x}$ with $[H_2]\bar{v} = [H_2']\bar{x}$ and all names in $\text{dom}(H_2') - \text{dom}$

are fresh.

Proof. Using the $(x_1, \dots, \square, \dots, v_n)$ context, view each value v individually. By induction on v . Unlike in stores, our values now also include lambdas:

Case $\lambda^{\bar{z}} \bar{x}. e$:

$$\begin{aligned} \Delta \mid \bar{z} \vdash \lambda^{\bar{z}} \bar{x}. e & \quad (1), \text{ by LAM} \\ \bar{z} = \text{roots}(H_2) & \quad (2), \text{ by invariant} \\ H_1 \otimes H_2 \mid \lambda^{\bar{z}} \bar{x}. e \longrightarrow_h H_1 \otimes (H_2, x \mapsto \lambda^{\bar{z}} \bar{x}. e) \mid x & \quad (3), (\text{lam}_h), \text{ fresh } x \\ [H_2, x \mapsto \lambda^{\bar{z}} \bar{x}. e]x = [H_2](\lambda^{\bar{z}} \bar{x}. e) & \quad (4), \text{ since } x \text{ fresh} \end{aligned}$$

Lemma 69. (*Dropping can progress*)

If $I((\text{drop } x; e), \Delta, \Gamma, H_1, H_2)$ then $H_1 \otimes H_2 \mid \text{drop } x; e \longrightarrow_h^* H_1 \otimes H_2' \mid e$ with $[H_1 \otimes H_2]e = [H_1 \otimes H_2']e$.

Proof.

$$\begin{aligned} \Delta \mid \Gamma, x \vdash \text{drop } x; e & \quad (1), \text{ by invariant} \\ \Delta \mid \Gamma \vdash e & \quad (2), \text{ by DROP} \\ \text{fv}(e) \in \Delta, \Gamma & \quad (3), \text{ by lemma 65} \\ \Gamma, x = \text{roots}(H_2) & \quad (4), \text{ by invariant} \\ (H_1 \otimes H_2)[x] = H_2[x] & \quad (5), \text{ by (4)} \\ H_1 \otimes H_2 \mid \text{drop } x; e \longrightarrow_h^* [(H_1 \otimes H_2)[x], H_1 \otimes H_2] \mid e & \quad (6), \text{ by lemma 62} \\ [(H_1 \otimes H_2)[x], H_1 \otimes H_2] = [H_2[x], H_1 \otimes H_2] & \quad (7), \text{ by (5)} \\ = [H_2[x], H_2] \otimes H_1 = H_1 \otimes [H_2[x], H_2] & \quad (8), \text{ by lemma 60} \\ [H_1 \otimes H_2]e = [H_2 \otimes [H_2[x], H_2]]e & \quad (9), \text{ by (3) and soundness} \end{aligned}$$

Lemma 70. (*Dropping can progress*)

If $I((\text{drop } x; e), \Delta, \Gamma, H_1, H_2)$ then $H_1 \otimes H_2 \mid \text{drop } x; e \longrightarrow_h^* H_1 \otimes H_2' \mid e$ with $[H_1 \otimes H_2]e = [H_1 \otimes H_2']e$.

Proof. Like proof of lemma 70, using lemma 63.

Lemma 71. (*Heap semantics can progress (no eval ctx)*)

If $I(e, \Delta, \Gamma, H_1, H_2)$ and $[H_1 \otimes H_2]e \longrightarrow e'$, then $H_1 \otimes H_2 \mid e \longrightarrow_h^* H_1 \otimes H_2' \mid e''$ with $e' = [H_1 \otimes H_2']e''$.

Proof. By case-analysis on $[H_1 \otimes H_2]e \longrightarrow e'$. We omit the cases which are unchanged from the store semantics.

Case (*dup*).

$$\begin{aligned} H_1 \otimes H_2 \mid \text{dup } x; e & \longrightarrow_h H_1 \otimes H_2 \otimes (H_1 \otimes H_2)[x] \mid e \quad (1), \text{ by lemma 61} \\ \text{dom}((H_1 \otimes H_2)[x]) \subseteq \text{dom}(H_1 \otimes H_2) & \quad (2), \text{ by definition} \\ [H_1 \otimes H_2]e = [H_1 \otimes H_2 \otimes (H_1 \otimes H_2)[x]]e & \quad (3), \text{ by (2)} \end{aligned}$$

Case (*drop*).

$$\begin{aligned} H_1 \otimes H_2 \mid \text{drop } x; e & \longrightarrow_h^* H_1 \otimes H_2' \mid e \quad (1), \text{ by lemma 69} \\ [H_1 \otimes H_2']e = [H_1 \otimes H_2]e & \quad (2), \text{ by lemma 69} \end{aligned}$$

Case (*dropru*).

$$\begin{aligned} H_1 \otimes H_2 \mid \text{dropru } x; e & \longrightarrow_h^* H_1 \otimes H_2' \mid e \quad (1), \text{ by lemma 70} \\ [H_1 \otimes H_2']e = [H_1 \otimes H_2]e & \quad (2), \text{ by lemma 70} \end{aligned}$$

Case (*alloc*).

$$\begin{aligned} H_1 \otimes H_2 \mid \text{alloc } k; e & \longrightarrow_h H_1 \otimes (H_2, \diamond_k) \mid e \quad (1), \text{ by } (\text{alloc}_h) \\ [H_1 \otimes H_2]e = [H_1 \otimes (H_2, \diamond_k)]e & \quad (2), \text{ clear} \end{aligned}$$

Case (beta).

$H_1 \otimes H_2 \mid (y) \bar{y} \longrightarrow_h H_1 \otimes H_2 \mid \text{dup } \bar{z}; \text{ drop } y; e[\bar{x}:=\bar{y}]$	(1), by (app_h), $(\lambda^{\bar{z}} \bar{x}. e) \in H_1 \otimes H_2$
$y, \bar{y} = \text{roots}(H_2)$	(2), by invariant
$(H_1 \otimes H_2)[\bar{z}] = H_2[\bar{z}]$ and $(H_1 \otimes H_2)[y] = H_2[y]$	(3), by (2)
$H_1 \otimes H_2 \mid (y) \bar{y} \longrightarrow_h H_1 \otimes [H_2[y], H_2 \otimes H_2[\bar{z}]] \mid e[\bar{x}:=\bar{y}]$	(4), by (3)
$\text{roots}([H_2[y], H_2 \otimes H_2[\bar{z}]]) = \bar{y}, \bar{z}$	(5), by (2)
$\text{fv}(e) = \bar{y}, \bar{z}$	(6), by lemma 65
$e' = [H_1 \otimes H_2']e''$	(7), by (5) and (6)

E.4 Heap semantics preserves linearity and roots

In this section, we wish to maintain the invariant $I(e, \Delta, \Gamma, H_1, H_2)$ which makes lemma 71 work.

Lemma 72. (*Variable substitution preserves FIP typing*)

Assuming that \bar{y} does not occur in e . (1) If $\Delta \mid \Gamma, \bar{x} \vdash e$, then $\Delta \mid \Gamma, \bar{y} \vdash e[\bar{x}:=\bar{y}]$. (2) If $\Delta, \bar{x} \mid \Gamma \vdash e$, then $\Delta, \bar{y} \mid \Gamma \vdash e[\bar{x}:=\bar{y}]$.

Proof. By induction on the FIP judgement for any such \bar{x}, \bar{y} .

Case LAM:

$\Delta \mid \bar{z}, \bar{x} \vdash \lambda^{\bar{z}, \bar{x}} \bar{x}'. e$	(1), given
$\emptyset \mid \bar{z}, \bar{x}, \bar{x}' \vdash e$	(2), by LAM
$\bar{z}, \bar{x} = \text{fv}(\lambda \bar{x}'. e)$	(3), by LAM
$\emptyset \mid \bar{z}, \bar{y}, \bar{x}' \vdash e[\bar{y}:=\bar{x}]$	(4), by LAM
$\bar{z}, \bar{y} = \text{fv}(\lambda \bar{x}'. e[\bar{y}:=\bar{x}])$	(5), by (3)
$\Delta \mid \bar{z}, \bar{y} \vdash \lambda^{\bar{z}, \bar{y}} \bar{x}'. e[\bar{y}:=\bar{x}]$	(6), by (4),(5)

Other cases are clear.

Lemma 73. (*The delta environment can be weakened*)

If $\Delta \mid \Gamma \vdash e$, then $\Delta, x \mid \Gamma \vdash e$.

Proof. By straight-forward induction on the judgement $\Delta \mid \Gamma \vdash e$. Unlike as in the proof of store semantics, we do not have to require $x \notin \Gamma$ as Δ and Γ need not be disjoint.

The next lemma is the main lemma of our soundness proof. It says that the individual steps of the heap semantics take sound/linear heaps to sound/linear heaps. As usual, the judgement \longrightarrow_h does not include the eval rule. We assume a slightly modified (app_h) rule, where we execute the dups immediately. This is equivalent to the (app_h) rule given in the paper:

$$(app_h) \quad H, \bar{z} \mapsto^n \bar{v} \mid (y) \bar{y} \longrightarrow_h \quad H, \bar{z} \mapsto^{n+1} \bar{v} \mid \text{drop } y; e[\bar{x}:=\bar{y}] \quad (y \mapsto^m \lambda^{\bar{z}} \bar{x}. e \in H)$$

Lemma 74. (*Heap semantics preserves linearity and roots (no eval ctx)*)

If $I(e, \Delta, \Gamma, H_1, H_2)$ and $H_1 \otimes H_2 \mid e \longrightarrow_h H_1 \otimes H_3 \mid e'$, then $I(e', \Delta', \Gamma', H_1, H_3)$.

Proof. By case analysis on the rules of the heap semantics.

Case (app_h).

FP²: Fully in-Place Functional Programming

- $\Delta \mid y, \bar{y} \vdash (y) \bar{y}$ (1), by APP
- $\emptyset \mid \bar{x}, \bar{z} \vdash e$ (2), by LAM
- $\emptyset \mid \bar{y}, \bar{z} \vdash e[\bar{x} := \bar{y}]$ (3), by 72
- $\emptyset \mid y, \bar{y}, \bar{z} \vdash \text{drop } y; e[\bar{x} := \bar{y}]$ (4), by DROP
- $H'_2, \bar{z} \mapsto^n \bar{v} := H_2$ (5), define
- $H_3 := H'_2, \bar{z} \mapsto^{n+1} \bar{v}$ (6), define
- $y, \bar{y} = \text{roots}(H_2)$ (7), given
- $y, \bar{y}, \bar{z} = \text{roots}(H_3)$ (8), by (7) and (6)

Case (*alloc_h*).

- $\Delta \mid \Gamma \vdash \text{alloc } k; e$ (1), given
- $H_3 := H_2, \diamond_k$ (2), define
- $\Delta \mid \Gamma, \diamond_k \vdash e$ (3), by ALLOC
- $\Gamma = \text{roots}(H_2)$ (4), by (1)
- $\Gamma, \diamond_k = \text{roots}(H_3)$ (5), by (2)
- H_3 linear (6), by (5)

Case (*lam_h*).

- $\Delta \mid \bar{z} \vdash \lambda \bar{x}. e$ (1), given
- $H_3 := H_2, x \mapsto \lambda \bar{x}. e$ (2), define
- $\Delta \mid x \vdash x$ (3), by VAR
- $x = \text{roots}(H_3)$ (4), by (2) and since x is fresh
- $\bar{z} = \text{roots}(H_2)$ (5), by (1)
- H_3 linear (6), by (5)

Case (*dcon_h*).

- $\Delta \mid \Gamma, x \vdash \text{drop } x; e$ (1), given
- $H_3 := H_2 - x$ (2), define
- $x \mapsto^1 C^k \bar{y} \in \text{roots}(H_2)$ (3), given
- H_3 sound (4), by (3)
- $\Gamma, \bar{y} = \text{roots}(H_3)$ (5), since H_2 is linear
- $\Delta \mid \Gamma, \bar{y} \vdash \text{drop } \bar{y}; e$ (6), by (1), DROP

Case (*dconru_h*).

- $\Delta \mid \Gamma, x \vdash \text{dropru } x; e$ (1), given
- $H_3 := H_2 - x, \diamond_k$ (2), define
- $x \mapsto^1 C^k \bar{y} \in \text{roots}(H_2)$ (3), given
- H_3 sound (4), by (3)
- $\Gamma, \diamond_k, \bar{y} = \text{roots}(H_3)$ (5), since H_2 is linear
- $\Delta \mid \Gamma, \diamond_k, \bar{y} \vdash \text{drop } \bar{y}; e$ (6), by (1), DROP

Case (*d_{lam}*).

- $\Delta \mid \Gamma, x \vdash \text{drop } x; e$ (1), given
- $H_3 := H_2 - x$ (2), define
- $x \mapsto^1 \lambda \bar{x}. e \in \text{roots}(H_2)$ (3), given
- H_3 sound (4), by (3)
- $\Gamma, \bar{y} = \text{roots}(H_3)$ (5), since H_2 is linear
- $\Delta \mid \Gamma, \bar{y} \vdash \text{drop } \bar{z}; e$ (6), by (1), DROP

Case (*drop_h*).

$\Gamma, x = \text{roots}(H_2)$ (1), given
 $H'_2, x \mapsto^{n+1} v := H_2$ (2), define
 $H_3 := H'_2, x \mapsto^n v$ (3), define
 $\Delta \mid \Gamma, x \vdash \text{drop } x; e$ (4), given
 H_3 sound (5), by (3)
 $\Gamma = \text{roots}(H_3)$ (6), since H_2 is linear
 $\Delta \mid \Gamma \vdash e$ (7), by (4)

Case (drop_h). Exactly as case (drop_h).

Case (dup_h).

$\Delta \mid \Gamma \vdash \text{dup } x; e$ (1), given
 $x \in (\Delta, \Gamma)$ (2), by (1)
 $H_3 := H_2 \otimes (H_1 \otimes H_2)[x]$ (3), by lemma 61
 $x \in \text{dom}(H_1 \otimes H_2)$ (4), by (2)
 H_3 linear (5), by (3),(4)
 $\Gamma, x = \text{roots}(H_3)$ (6), where $\Gamma = \text{roots}(H_2)$
 $\Delta \mid \Gamma, x \vdash e$ (7), by (1)

E.5 Main Invariant

As before, we want to generalize lemma 74 to arbitrary evaluation contexts and use the more complicated invariant. The simple invariant fails here for the same reason it failed in the store semantics – and all results from that section transfer directly to the heap semantics. This is due to the fact that the heap and store semantics use the same evaluation context and borrowing strategy (in the let-bindings).

We define $I_E(e, \Delta, \Gamma, H_1, H_2)$ by induction on E :

Case $E = \square$:

$I_{\square}(e, \Delta, \Gamma, H_1, H_2) := I(e, \Delta, \Gamma, H_1, H_2)$

Case $E[E'] = C^k x_1 \dots E' \dots v_k \mid (x_1, \dots, E', \dots, v_n) \mid E' e \mid x E' \mid f(\bar{y}; E') \mid \text{match } E' \{ \overline{p \mapsto e} \}$:

$I_{E[E']}(e, \Delta, (\Gamma_i, \Gamma), H_1, H_2) := I_{E'}(e, \Delta', \Gamma_i, (H_1 \otimes [H_2[\Gamma_i], H_2]), H_2[\Gamma_i])$ and $I(E[\square], \Delta, \Gamma, H_1, [H_2[\Gamma_i], H_2])$

Case $E[E'] = \text{let } \bar{x} = E' \text{ in } e$:

$I_{E[E']}(e, \Delta, (\Gamma_1, \Gamma_2, \Gamma_3), H_1, H_2) :=$
 $I_{E'}(e, \Delta', \Gamma_1, (H_1 \otimes [H_2[\Gamma_i], H_2]), H_2[\Gamma_i])$
 and $I(E[\square], \Delta, (\Gamma_2, \Gamma_3), H_1, [H_2[\Gamma_i], H_2])$

All of lemmas about exchanging the two invariants continue to hold largely unchanged. The only difference is that we have to join $H_2[\Gamma_i]$ onto $[H_2[\Gamma_i], H_2]$ using lemma 57, which guarantees that this yields the original H_2 .

Lemma 75. (*Weakening the invariant*)

If $I(E[e], \Delta, \Gamma, H_1, H_2)$, then $I_E(e, \Delta, \Gamma, H_1, H_2)$.

Lemma 76. (*Weakening the invariant*)

If $I_E(E_2[e], \Delta, \Gamma, H_1, H_2)$, then $I_{E[E_2]}(e, \Delta, \Gamma, H_1, H_2)$.

Lemma 77. (*Values do not borrow*)

If $\Delta \mid \Gamma \vdash \bar{v}$, then $\emptyset \mid \Gamma \vdash \bar{v}$.

Proof. By induction on $\Delta \mid \Gamma \vdash \bar{v}$:

Case $\lambda \bar{x}. \bar{e}$:

FP²: Fully in-Place Functional Programming

- $\Delta \mid \bar{z} \vdash \lambda \bar{x}. e \quad (1), \text{ by LAM}$
- $\emptyset \mid \bar{z}, \bar{x} \vdash e \quad (2), \text{ by (1)}$
- $\emptyset \mid \bar{z} \vdash \lambda \bar{x}. e \quad (3), \text{ by LAM}$

Lemma 78. (*Values do not borrow*)

If $I(\bar{v}, \Delta, \Gamma, H_1, H_2)$, then $I(\bar{v}, \emptyset, \Gamma, H'_1, H_2)$ for any H'_1 .

However, notice that we introduce a Δ' enviroment at every level of the evaluation context. That means, that we can only remove a single level of the evaluation context at a time. To model this, we call an evaluation context E flat if $E = E'[E'']$ implies that either E' or E'' is the hole. Then:

Lemma 79. (*Strengthening the invariant*)

If E is flat and $I_E(\bar{v}, \Delta, \Gamma, H_1, H_2)$, then $I(E[\bar{v}], \Delta, \Gamma, H_1, H_2)$.

Lemma 80. (*Strengthening the invariant*)

If E_2 is flat and $I_{E_1}(\bar{v}, \Delta, \Gamma, H_1, H_2)$, then $I_E(E_2[\bar{v}], \Delta, \Gamma, H_1, H_2)$.

E.6 Soundness

Now we can show the main soundness theorem. As earlier, we extend the progress and preservation proofs to handle the `STEP` and `eval` cases. If we evaluate e_1 under the context E_1 , we need to assume the invariant $I_{E_1}(e_1, \Delta, \Gamma, H_1, H_2)$. But how can we obtain the invariant for this precise E_1 ? The trick is that we do not have to know E_1 , instead we can just assume $I_{E_2}(e_2, \Delta, \Gamma, H_1, H_2)$ for $E_1[e_1] = E_2[e_2]$:

Lemma 81. (*Comparing evaluation contexts*)

Let $E_1[e_1] = E_2[e_2]$, then either:

- $E_1 = E'_1[E''_1]$ with $E'_1 = E_2$ and $e_2 = E''_1[e_1]$
- $E_2 = E'_2[E''_2]$ with $E'_2 = E_1$ and $e_1 = E''_2[e_2]$

Lemma 82. (*Alignment of invariant*)

If $e_1 \longrightarrow e'_1$, $E_1[e_1] = E_2[e_2]$ and $I_{E_2}(e_2, \Delta, \Gamma, H_1, H_2)$, then $I_{E_1}(e_1, \Delta, \Gamma, H_1, H_2)$.

Lemma 83. (*Heap semantics can progress*)

If $I_E(e, \Delta, \Gamma, H_1, H_2)$ and $[H_1 \otimes H_2]e \longrightarrow e'$, then $H_1 \otimes H_2 \mid e \longrightarrow_h^* H_1 \otimes H'_2 \mid e''$ with $e' = [H_1 \otimes H'_2]e''$.

Proof.

- $I_E(e, \Delta, \Gamma, H_1, H_2) \quad (1), \text{ given}$
- $I(e, \Delta', \Gamma', H'_1, H'_2), \text{ with } H_1 \otimes H_2 = H'_1 \otimes H'_2 \text{ and } H'_2 \subseteq H_2 \quad (2), \text{ by (1)}$
- $[H_1 \otimes H_2]e = [H'_1 \otimes H'_2]e \quad (3), \text{ by (2)}$
- $H'_1 \otimes H'_2 \mid e \longrightarrow_h^* H'_1 \otimes H'_3 \mid e'' \quad (4), \text{ by lemma 71}$
- $e' = [H'_1 \otimes H'_3]e'' \quad (5), \text{ by lemma 71}$
- $H_3 := H'_3 \otimes [H'_2, H_2] \quad (6), \text{ define}$
- $H_1 \otimes H_2 \mid e \longrightarrow_h^* H_1 \otimes H_3 \mid e'' \quad (7), \text{ by (4)}$
- $e' = [H_1, H_3]e'' \quad (8), \text{ by (5)}$

Lemma 84. (*Heap semantics preserves linearity and roots*)

If $I_E(e, \Delta, \Gamma, H_1, H_2)$ and $H_1 \otimes H_2 \mid e \longrightarrow_h H_1 \otimes H_3 \mid e'$, then $I_E(e', \Delta', \Gamma', H_1, H_3)$.

Proof.

- $I_E(e, \Delta, (\Gamma, \Gamma'), H_1, (H_2 \otimes H'_2))$ (1), given
 $I(e, \Delta', \Gamma, (H_1 \otimes H'_2), H_2)$ (2), split (1)
 $I_E(\cdot, \Delta, \Gamma', H_1, H'_2)$ (3), split (1)
 $H_1 \otimes H'_2 \otimes H_2 \mid e \longrightarrow_h H_1 \otimes H'_2 \otimes H'_3 \mid e'$ (4), by assumption
 $I(e', \Delta'', \Gamma'', (H_1 \otimes H'_2), H'_3)$ (5), by lemma 74
 $H_3 := H'_3 \otimes H'_2$ (6), define
 $I_E(e', \Delta, (\Gamma'', \Gamma'), H_1, H_3)$ (7), merge (3) and (5)

Lemma 85. (*Soundness lemma*)

If $I_E(e, \Delta, \Gamma, H_1, H_2)$ and $[H_1 \otimes H_2](E[e]) \longmapsto^* \bar{v}$, then $H_1 \otimes H_2 \mid E[e] \longmapsto_h^* H_1 \otimes H_3 \mid \bar{x}$ with $I(\bar{x}, \emptyset, \bar{x}, \emptyset, H_3)$ and $[H_3]\bar{x} = \bar{v}$.

Proof. By induction on $[H_1, H_2](E[e]) \longmapsto^* \bar{v}$.

Case Reflexive case:

- $[H_1 \otimes H_2](E[e]) = \bar{v}$ (1), given
 $E = \square, e = \bar{w}, [H_1 \otimes H_2]\bar{w} = \bar{v}$ (2), by (1)
 $I_\square(\bar{w}, \Delta, \Gamma, H_1, H_2)$ (3), given
 $I(\bar{w}, \Delta, \Gamma, H_1, H_2)$ (4), by definition
 $H_1 \otimes H_2 \mid \bar{w} \longrightarrow_h^* H_1 \otimes H'_2 \mid \bar{x}$ (5), by lemma 68
 $\bar{v} = [H_1 \otimes H_2]\bar{w} = [H_1 \otimes H'_2]\bar{x}$ (6), by lemma 68
 $I(\bar{x}, \Delta, \Gamma, H_1, H'_2)$ (7), by lemma 74
 $I(\bar{x}, \emptyset, \Gamma, \emptyset, H'_2)$ (8), by lemma 78

Case Transitive case:

- $[H_1 \otimes H_2](E[e]) \longmapsto e'$ (1), given
 $e' \longmapsto^* \bar{v}$ (2), given
 $E'_1[e'_1] = [H_1 \otimes H_2](E[e]), e'_1 \longrightarrow e'_2, e' = E'_1[e'_2]$ (3), by STEP
 $[H_1 \otimes H_2]E_1 := E'_1, [H_1 \otimes H_2]e_1 := e'_1$ (4), define
 $E_1[e_1] = E[e]$ (5), by (4)
 $I_E(e, \Delta, \Gamma, H_1, H_2)$ (6), given
 $I_{E_1}(e_1, \Delta, \Gamma, H_1, H_2)$ (7), by lemma 82
 $H_1 \otimes H_2 \mid e_1 \longrightarrow_h^* H_1 \otimes H'_2 \mid e''$ (8), by lemma 83
 $e'_2 = [H_1 \otimes H'_2]e''$ (9), by lemma 83
 $I_{E_1}(e'', \Delta', \Gamma', H_1, H'_2)$ (10), by lemma 84
 $[H_1 \otimes H'_2](E_1[e'']) = ([H_1 \otimes H'_2]E_1)([H_1 \otimes H'_2]e'')$ (11), commute
 $= E'_1[e'_2]$ (12), by (4) and (9)
 $= e' \longmapsto^* \bar{v}$ (13), by (3) and (2)
 $H_1 \otimes H'_2 \mid E_1[e''] \longmapsto_h^* H_1 \otimes H_3 \mid \bar{x}$ (14), inductive hypothesis
 $I(\bar{x}, \emptyset, \bar{x}, \emptyset, H_3)$ and $[H_3]\bar{x} = \bar{v}$ (15), inductive hypothesis
 $H_1 \otimes H_2 \mid E_1[e_1] \longmapsto_h^* H_1 \otimes H'_2 \mid E_1[e'']$ (16), EVAL on (8)
 $H_1 \otimes H_2 \mid E[e] \longmapsto_h^* H_1 \otimes H_3 \mid \bar{x}$ (17), append (16) and (14)

Theorem 10. (*FIP programs are sound in heap semantics*)

If $\Delta \mid \Gamma \vdash e$ and given compatible heaps H_1, H_2 with $\Delta \subseteq \text{dom}(H_1)$, H_1 sound, $\Gamma = \text{roots}(H_2)$ and H_2 linear, and $[H_1 \otimes H_2]e \longrightarrow^* \bar{v}$, then $H_1 \otimes H_2 \mid e \longrightarrow_h^* H_1 \otimes H_3 \mid \bar{x}$ where $[H_3]\bar{x} = \bar{v}$, $\bar{x} = \text{roots}(H_3)$ and H_3 is linear.

Proof.

FP²: Fully in-Place Functional Programming

$\Delta \mid \Gamma \vdash e$	(1), given
H_1, H_2 compatible heaps	(2), given
$\Delta \subseteq \text{dom}(H_1)$ and H_1 sound	(3), given
$\Gamma = \text{roots}(H_2)$ and H_2 linear	(4), given
$I(e, \Delta, \Gamma, H_1, H_2)$	(5), by (1)-(4)
$e = E[e']$	(6), for some E, e'
$I_E(e', \Delta, \Gamma, H_1, H_2)$	(7), lemma 75
$H_1 \otimes H_2 \mid E[e'] \mapsto_h^* H_1 \otimes H_3 \mid \bar{x}$	(8), lemma 85
$I(\bar{x}, \emptyset, \bar{x}, \emptyset, H_3)$	(9), lemma 85
$[H_3]\bar{x} = \bar{v}$	(10), lemma 85
$\bar{x} = \text{roots}(H_3)$ and H_3 linear	(11), by (9)

Received 2023-07-07; accepted 2023-11-07